

FOREWORD

This publication is issued by the National Computer Security Center (NCSC) as part of its program to promulgate technical computer security guidelines. The interpretations extend the evaluation classes of the Trusted Systems Evaluation Criteria (DOD 5200.28-STD) to trusted network systems and components.

This document will be used for a period of at least one year after date of signature. During this period the NCSC will gain experience using the Trusted Network Interpretation in several network evaluations. In addition, the NCSC will conduct a series of tutorials and workshops to educate the community on the details of the Trusted Network Interpretation and receive feedback. After this trial period, necessary changes to the document will be made and a revised version issued.

Workshops and tutorials will be open to any member of the network security community interested in providing feedback. Anyone wishing more information, or wishing to provide comments on the usefulness or correctness of the Trusted Network Interpretation may contact: Chief, Technical Guidelines Division, National Computer Security Center, Ft. George G. Meade, MD 20755-6000, ATTN: C11. The telephone number is (301) 859-4452.

PATRICK R. GALLAGHER, JR. 31 July 1987
Director
National Computer Security Center

ACKNOWLEDGMENT

Acknowledgment is extended to the members of the Working Group who produced this Interpretation. Members were: Alfred Arsenault, National Computer Security Center (Chair); Dr. Roger Schell, Gemini Computers; Stephen Walker, Trusted Information Systems; Dr. Marshall Abrams, MITRE; Dr. Jonathan Millen, MITRE; Leonard LaPadula, MITRE; Robert Morris, NCSC; and Jack Moskowitz, NCSC. Also due acknowledgement for their many inputs to this interpretation are Steve Padilla and William Shockley, Gemini Computers.

Introduction

I.1. Scope

Part I of this document provides interpretations of the Department of Defense Trusted Computer System Evaluation Criteria (TCSEC) (DOD-5200.28-STD), for trusted computer/communications network systems. The specific security feature, the assurance requirements, and the rating structure of the TCSEC are extended to networks of computers ranging from isolated local area networks to wide-area internetwork systems.

Part II of this document describes a number of additional security services (e.g., communications integrity, denial of service, transmission security) that arise in conjunction with networks. Those services available in specific network offerings, while inappropriate for the rigorous evaluation applied to TCSEC related feature and assurance requirements, may receive qualitative ratings.

The TCSEC related feature and assurance requirements, and the additional security services described herein are intended for the evaluation of trusted network systems designed to protect a range of sensitive information. As such, they require that physical, administrative, procedural, and related protection measures adequate to the sensitivity of the information being handled are already in place. It is possible, and indeed common practice, to operate a network in a secure manner at a single system high sensitivity level without meeting any trust related feature or assurance requirements described herein. The full range of physical and administrative security measures appropriate to the highest sensitivity level of information on the network must be in place in order to operate a trusted network as described in this Interpretation.

It is important to note that this Interpretation does not describe all the security requirements that may be imposed on a network. Depending upon the particular environment, there may be communications security (COMSEC), emanations security, physical security, and other measures required.

An environmental evaluation process, such as that described in the ``Computer Security Requirements--Guidance for Applying the DoD TCSEC in Specific Environments'' (CSC-

STD-003-85), can be used to determine the level of trust required by specific system environments. Similar analyses are applicable to networks evaluated under these Interpretations.

I.2. Purpose

As with the TCSEC itself, this Interpretation has been prepared for the following purposes:

1. To provide a standard to manufacturers as to what security features and assurance levels to build into their new and planned, commercial network products in order to provide widely available systems that satisfy trust requirements for sensitive applications
2. To provide a metric by which to evaluate the degree of trust that can be placed in a given network system for processing sensitive information
3. To provide a basis for specifying security requirements in acquisition specifications.

With respect to the second purpose for development of the criteria, i.e., providing a security evaluation metric, evaluations can be delineated into two types: an evaluation can be performed on a network product from a perspective that excludes the application environment, or an evaluation can be done to assess whether appropriate security measures have been taken to permit the system to be used operationally in a specific environment. The former type of evaluation is done by the National Computer Security Center through the Commercial Product Evaluation Process.

The latter type of evaluation, those done for the purpose of assessing a system's security attributes with respect to a specific operational mission, is known as a certification evaluation. It must be understood that the completion of a formal product evaluation does not constitute certification or accreditation for the system to be used in any specific application environment. On the contrary, the evaluation report only provides a trusted network system's evaluation rating along with supporting data describing the product system's strengths and weaknesses from a computer security point of view. The system security certification and the formal approval/accreditation procedure, done in accordance with the applicable policies of the issuing agencies, must still be followed before a network can be approved for use in processing or handling classified information. Designated Approving Authorities (DAAs)

remain ultimately responsible for specifying the security of systems they accredit.

This Interpretation can be used directly and indirectly in the certification process. Along with applicable policy, it can be used directly as technical guidance for evaluation of the total system and for specifying system security and certification requirements for new acquisitions. Where a system being evaluated for certification employs a product that has undergone a Commercial Product Evaluation, reports from that process will be used as input to the certification evaluation. Technical data will be furnished to designers, evaluators, and the DAAs to support their needs for making decisions.

The fundamental computer security requirements as defined in the TCSEC apply to this Interpretation.

I.3. Background

The term ``sponsor'' is used throughout this document to represent the individual or entity responsible for presenting a component or network system for evaluation. The sponsor may be a manufacturer, vendor, architect, developer, program manager, or related entity.

A network system is the entire collection of hardware, firmware, and software necessary to provide a desired functionality.

A component is any part of a system that, taken by itself, provides all or a portion of the total functionality required of a system. A component is recursively defined to be an individual unit, not useful to further subdivide, or a collection of components up to and including the entire system.

Because the term integrity has been used in various contexts to denote specific aspects of an overall issue, it is important for the reader to understand the context in which the term is used in this document. Within Part I, as in the TCSEC itself, the use of this term is limited to (1) the correct operation of NTCB hardware/firmware and (2) protection against unauthorized modification of labels and data. Specifically, all NTCBs that support sensitivity labels (viz., Divisions A and B) must, as detailed in the Label Integrity section of the TCSEC, protect the labels that represent the sensitivity of information (contained in objects) and the corresponding authorizations of users (with subjects as surrogates). In addition, for network systems with a defined data integrity policy, the NTCB must control the accesses of users that modify information-. As part of the NTCB itself, such integrity policies will be supported

by access control mechanisms based on the identity of individuals (for discretionary integrity) and/or sensitivity levels (for mandatory integrity). In contrast, within Part II the term integrity relates to the mechanisms for information transfer between distinct components. This communications integrity includes the issues for correctness of message transmission, authentication of source/destination, data/control/protocol communication field correctness and related areas.

- See, for example, K. J. Biba, Integrity Considerations for Secure Computer Systems, MTR-3153, The MITRE Corporation, Bedford, MA, June 1975.

In many network environments, encryption can play an essential role in protecting sensitive information. In Part I of this document, encryption is referenced as a basis for providing data and label integrity assurance. In Part II, encryption is described as a tool for protecting data from compromise or modification attacks. Encryption algorithms and their implementation are outside the scope of this document. This document was prepared from a DoD perspective and requires NSA approval relative to the use of encryption. In other contexts, alternate approval authority may exist.

As with the TCSEC itself, this is a reference document and is not intended to be used as a tutorial. The reader is assumed to be familiar with the background literature on computer security and communications networks. Part II assumes a familiarity with the terminology used within ISO Security Services documentation.

= See, for example, M. D. Abrams and H. J. Podell, Tutorial: Computer and Network Security, IEEE Computer Society Press, 1987.

* ISO 7498/Part 2 - Security Architecture, ISO / TC 97 / SC 21 / N1528 / WG 1 Ad hoc group on Security, Project 97.21.18, September 1986.

I.3.1. Trusted Computer System Evaluation Criteria

The DoD TCSEC was published in December 1985 to provide a means of evaluating specific security features and

assurance requirements available in ``trusted commercially available automatic data processing (ADP) systems,'' referred to in this document as Automated Information System (AIS). The rating scale of the TCSEC extends from a rating which represents a minimally useful level of trust to one for ``state of the art'' features and assurance measures. These technical criteria guide system builders and evaluators in determining the level of trust required for specific systems. When combined with environmental guidelines, minimum security protection requirements, and appropriate accreditation decisions for specific installations can be determined. The philosophy of protection embodied in the TCSEC requires that the access of subjects (i.e., human users or processes acting on their behalf) to objects (i.e., containers of sensitive information) be mediated in accordance with an explicit and well defined security policy.

In order to ensure strict compatibility between TCSEC evaluated AIS and networks and their components, and to avoid the possible evolution of incompatible evaluation criteria, Part I of this document has been specifically prepared as an INTERPRETATION of the TCSEC for networks. It is based entirely on the principles of the TCSEC, uses all TCSEC basic definitions, and introduces new concepts only where essential to understanding the TCSEC in a network context. Unless otherwise stated, the TCSEC requirements apply as written. The approach of interpreting the TCSEC for networks in general has already been successfully employed in a number of specific complex network and AIS applications.

There are several security policy models that may be used with the reference monitor concept. The Bell-LaPadula model is commonly used but is not mandated. Similarly for integrity policy, models such as Biba have been proposed but are not mandated. For this network interpretation, no specific access control policy is required; however, it is necessary that either a secrecy policy, an integrity policy, or both be specified for enforcement by the reference monitor.

In the context of network systems, there are a number of additional security services that do not normally arise in individual AIS, and are not appropriate to the detailed feature and assurance evaluation prescribed by the TCSEC. These security services, which may or may not be available in specific network offerings, include provisions for communications security, denial of service, transmission security, and supportive primitives, such as encryption mechanisms and protocols. Part II of this document describes these services and provides a qualitative means of evaluating their effectiveness when provided.

Evaluation of Part II offered services is dependent upon the results of the system's Part I evaluation or

component's Appendix A evaluation. A Part II evaluation rating of good in a component or system which has a low Part I trust rating is of limited value. The sponsor must identify which security services are offered by a system or component for evaluation against Part II. The evaluators will normally give a none, minimum, fair or good rating for those services offered. In some cases, a rating of present is all that can be given or a quantitative measure of strength may be used as the basis for rating. A not applicable rating will be given for services not offered by the product. Part II services may be provided by mechanisms outside the NTCB.

I.3.2. Two Network Views

DoD Directive 5200.28 (and similar policies elsewhere in government) establishes the concept of a DAA, an individual who is responsible for approving the use of an AIS for processing classified information. For stand-alone AIS, this approval process and the technical advisory role to the DAA provided by the TCSEC are well understood. The same approval process applies to networks of AIS and plays a key role in determining how and when networks can be evaluated using this Interpretation.

Depending upon the operational and technical characteristics of the environment in which a network exists, either of two views for accreditation and evaluation purposes applies: as a collection of two or more interconnected separately accredited AIS or as a single unified system the security accreditation of which is the responsibility of a single authority.

The security feature and assurance requirements of a specified network, and hence its suitability for evaluation under this Interpretation, is greatly impacted by which view of the network is appropriate.

I.3.2.1. Interconnected Accredited AIS View

The interconnected accredited AIS view is an operational perspective that recognizes that parts of the network may be independently created, managed, and accredited. Where different accrediting jurisdictions are involved, the joint approval process is required, describing the handling practices and classification levels that will be exchanged between the components involved.

Interconnected accredited AIS consist of multiple systems (some of which may be trusted) that have been independently assigned operational sensitivity ranges (the highest and lowest sensitivity levels of information that may be

simultaneously processed on that system). In this view, the individual AIS may be thought of as ``devices'' with which neighboring systems can send and receive information. Each AIS is accredited to handle sensitive information at a single level or over a range between a minimum and maximum level.

The range of sensitive information that may be exchanged between two such AIS is a range, agreed upon by each system's approving authorities, which cannot exceed the maximum sensitivity levels in common between the two systems.

Because of the complex structure of a network consisting of interconnected accredited AIS, it may not be practical to evaluate such a network using this Interpretation or to assign it a trusted system rating. In this case, the accreditor is forced to accept the risk of assessing the security of the network without the benefit of an evaluation against the principles of the TCSEC as interpreted in Part I of the document. Appendix C describes the rules for connecting separately accredited AIS and the circumstances in which these rules apply.

I.3.2.2. Single Trusted System View

The policy enforcement by trusted components in a ``single trusted system'' exhibits a common level of trust throughout. A single trusted system is accredited as a single entity by a single accrediting authority. (In certain circumstances where a system will process information from multiple sensitive sources, more than one accrediting authority may be involved, but their responsibility will be for accrediting the whole system as a single entity for use processing the information for which they have authority.) Networks such as these can be evaluated against this Interpretation and given a rating compatible with trusted AIS evaluated by the TCSEC itself.

A ``single trusted system'' network implements a reference monitor to enforce the access of subjects to objects in accordance with an explicit and well defined network security policy. The network has a single trusted computing base, referred to as the Network Trusted Computing Base (NTCB), which is partitioned (see section I.4.2) among the network components in a manner that ensures the overall network security policy is enforced by the network as a whole.

Every component that is trusted must enforce a component-level security policy that may contain elements of the overall network security policy. The sum of all component-level security policies must be shown to enforce the overall network security policy.

There is no requirement that every component in the network have an NTCB partition nor that any such partition comprise a complete TCB (e.g., a network component could be dedicated to supporting the audit function and implement only that portion of the NTCB). Interaction among NTCB partitions shall be via communications channels, operating at either single or multiple levels as appropriate. The network security architect must identify how the NTCB is partitioned and how all the trusted system requirements are satisfied.

A given component that does not enforce the full implementation of all polices (i.e., mandatory access control, discretionary access control, identification/ authentication and audit) must be evaluated as a component as specified in Appendix A. For example, a network architecture that does not operate above Level 3 of the ISO protocol model and typically does not enforce discretionary access control must be evaluated as a component under Appendix A and not as a full system.

I.3.2.2.1. Connection-Oriented Abstraction

In many networking environments, the overall network security policy includes controlling the establishment of authorized connections across the network. The access control mediation performed by the components of these networks enforces the establishment of connections between host computers on the network in accordance with some form of authorized connection list. While a connection-oriented policy may be suitable from an overall network perspective, specifying such a policy in terms of component level abstractions may be difficult but is required in order to evaluate the network.

Individual trusted network components may employ a local mechanism to enforce mediation only between local subjects and objects, as described in the TCSEC. Some of these components may have no direct involvement with the enforcement of network connections. Others, however, will have an additional higher level network connection enforcement role. This higher level connection-oriented abstraction may be enforced solely within an individual component or may be distributed across many components (e.g., in the end-to-end encryption case, cryptographic front end devices enforce the network connection authorization decisions made by an access control/key management center.)

With the connection-oriented abstraction, the role of the network as a whole in controlling information flow may be more easily understood, but there may be no simple way to extend this abstraction to the reference monitor require-

ments of individual components in the network. The overall network security policy must be decomposed into policy elements that are allocated to appropriate components and used as the basis for security policy models for these components.

The reference monitor subject/object definitions as given in the TCSEC represent the fundamental security policy enforcement at the individual component level but may not directly describe the overall network security policy issues such as the network's connection policy. The connection-oriented abstraction may be essential to understanding the overall network security policy. The network architecture must demonstrate the linkage between the connection-oriented abstraction and its realization in the individual components of the network.

I.3.2.2.2. Subjects and Objects

For purposes of this trusted network interpretation, the terms ``subject'' and ``object'' are defined as in the TCSEC.

The subjects of a trusted network commonly fall into two classes: those that serve as direct surrogates for a user (where ``user'' is synonymous with ``human being''), and ``internal'' subjects that provide services for other subjects--typically representing software process rather than being made part of each user surrogate subject.

There is a set of TCSEC requirements that are directed at users, rather than subjects. In the network context, services used to facilitate communications between users and AIS (e.g., protocol handlers) are usually provided by internal subjects. Some components that provide only communications facilitating services have only internal subjects.

Examples of ``single trusted system'' networks or components could include- packet-switched communications sub-networks (as found in the Defense Data Network (DDN), end-to-end (or host-to-host) encryption systems (such as used in Blacker or ANSI X9.17 implementations), application level networks or closed user community systems (such as the Inter Service/Agency Automated Message Processing Exchange [I S/A AMPE] and SACDIN Programs), local area networks, digital PABX systems, private switched networks (such as circuit-switched telecommunications systems), future Integrated Services Digital Network (ISDN) implementations, and a Virtual Machine Monitor (VMM) on a single computer when analyzed as a network.

I.4. Evaluation of Networks

The TCSEC provides a means for evaluating the trustworthiness of a system and assigning an evaluation class based on its technical properties - independent of the particular use for or the sensitivity of information being processed on the system. In this Interpretation, a network as a whole with its various interconnected components is recognized as a special instance of a trusted system. The designer of a trusted system is unconstrained by the TCSEC on design and implementation choices as long as for the

- Examples are employed throughout this document to clarify the concepts presented. The naming of an example implies no judgement of the product or system named nor on its suitability for any particular purpose.

system as a whole there is a clearly distinguished TCB with a definitive protection domain boundary. The features and assurance measures provided within the TCB perimeter will determine the evaluation class. The network must be viewed as PARTITIONED into a set of interconnected components, where each component may have an independent ``NTCB partition.'' All interaction between such trusted components must be via ``communication channels or I/O devices'' as defined by the TCSEC. For Division A and B networks these will generally be ``multilevel devices.''

I.4.1. Network Security Architecture and Design

Any network evaluated under this Interpretation must possess a coherent Network Security Architecture and Design. (Interconnection of components that do not adhere to such a Network Security Architecture is addressed in the Interconnection Rules, Appendix C.) The Network Security Architecture must address the security-relevant policies, objectives, and protocols. The Network Security Design specifies the interfaces and services that must be incorporated into the network so that it can be evaluated as a trusted entity. There may be multiple designs that conform to the same architecture but which are more or less incompatible and non-interoperable (except through the Interconnection Rules). Security related mechanisms that require cooperation among components are specified in the design in terms of their visible interfaces; mechanisms which have no visible interfaces are not specified in this document but are left as implementation decisions.

The Network Security Architecture and Design must be available from the network sponsor before evaluation of the network, or any component, can be undertaken. The Network Security Architecture and Design must be sufficiently com-

plete, unambiguous, and free from obvious flaws to permit the construction or assembly of a trusted network based on the structure it specifies.

When a component is being designed or presented for evaluation, or when a network assembled from components is assembled or presented for evaluation, there must be a priori evidence that the Network Security Architecture and Design are satisfied. That is, the components are assemblable into a network that conforms in every way with the Network Security Architecture and Design to produce a physical realization which is trusted to the extent that its evaluation indicates.

In order for a trusted network to be constructed from components that can be built independently, the Network Security Architecture and Design must completely and unambiguously define the security functionality of components as well as the interfaces between or among components. The Network Security Architecture and Design must be evaluated to determine that a network constructed to its specifications will in fact be trusted, that is, it will be evaluatable under these Interpretations.

I.4.2. The Partitioned NTCB

Like a stand-alone system, the network as a whole possesses a single TCB, referred to as the NTCB, consisting of the totality of security-relevant portions of the network. But, unlike a stand-alone system, the design and evaluation of the network rests on an understanding of how the security mechanisms are distributed and allocated to various components, in such a way that the security policy is supported reliably in spite of (1) the vulnerability of the communication paths and (2) the concurrent, asynchronous operation of the network components.

Some distributed systems have reliable, protected communication paths and thus satisfy only the first characteristic of a network: the division into concurrently operating, communicating processing components. Although certain interpretations in this Interpretation will not apply to them, it may be beneficial to employ this Interpretation to evaluate them, and to take advantage of the interpretations relating to component properties and interfaces.

An NTCB that is distributed over a number of network components is referred to as partitioned, and that part of the NTCB residing in a given component is referred to as an NTCB partition. A network host may possess a TCB that has previously been evaluated as a stand-alone system. Such a TCB does not necessarily coincide with the NTCB partition in

the host, in the sense of having the same security perimeter. Whether it does or not depends on whether the security policy for which the host TCB was evaluated coincides with the network security policy, to the extent that it is allocated to that host.

Even when a network host has a TCB that has been previously evaluated at a given class, and the host's TCB coincides with the host's NTCB partition, there is still no a priori relationship between the evaluation class of the host and the evaluation class of the network. Some examples will be given below to illustrate this point.

To evaluate a network at a given class, each requirement in Part I for that class must be satisfied by the network as a whole. It is also necessary to understand how each requirement is allocated among the network's components. Some components, such as the hosts, may satisfy the entire security policy in isolation; others, such as packet switches and access control centers, may have more specialized functions that satisfy only a subset of the network security policy. In addition, distinct subsets of the network security policy may be allocated to different network components.

Forcing every component to satisfy a specific Part I requirement is neither necessary nor sufficient to ensure that the network as a whole meets that requirement.

To show that it is not sufficient, consider two trusted multilevel AIS that export and import labeled information to and from each other over a direct connection. Both satisfy the Label Integrity requirement that a sensitivity label be accurately and unambiguously associated with exported data. If they were to have different, incompatible label encodings for the same sensitive information, the network as a whole would fail to satisfy the label integrity requirement. As a result, these Interpretations require at the B1 level and above that there be uniform labeling of sensitive information throughout the network.

To show that it is not necessary, consider the Mandatory Access Control requirement that at least two sensitivity levels be supported. Suppose that the network consists of a number of untrusted hosts that are incapable of maintaining labels and are operating at different levels in a single-level mode. If they are interconnected through suitable multilevel interface units, the network as a whole can support the "two or more levels" requirement.

The allocation of a requirement to a component does not simply mean that the component satisfies the requirement in isolation, but includes the possibility that it depends on other components to satisfy the requirement locally, or

cooperates with other components to ensure that the requirement is satisfied elsewhere in the network.

Taken together, these examples illustrate the essential role of an overall network security architecture in designing and evaluating a trusted network.

I.4.3. Component Evaluation

Because network components are often supplied by different vendors and are designed to support standardized or common functions in a variety of networks, significant advantages can accrue from a procedure for evaluating individual components. The purpose of component evaluation is to aid both the network designer and the evaluator by performing the evaluation process once and reusing the results whenever that component is incorporated into a network.

There are four types of security policies that may be supported by a network component:

1. Mandatory Access Control
2. Discretionary Access Control
3. Supportive policies (e.g., Authentication, Audit)
4. Application policies (e.g., the policy supported by a DBMS that is distinct from that supported by the underlying system)

Application level policies are user dependent and will not be considered further in these Interpretations.

For a component to support a policy such as Mandatory Access Controls, it must support all the required features for that policy with all of the required assurances of the given class.

I.5. Structure of the Document

The remainder of this document is divided into two parts, three appendices, a list of acronyms, a glossary, and a list of references. Part I presents TCSEC statements and detailed interpretations, which together constitute the requirements against which networks will be evaluated; and rationale for the network interpretation of the TCSEC. The TCSEC statement applies as modified by the Interpretation. Part II is a description of other Security Services not covered in the TCSEC interpretation which may be applicable to networks. Appendix A describes the evaluation of network components. Appendix B describes the rationale for network

partitioning into individual components. Appendix C describes the interconnect rules for linking interconnected accredited AIS.

Part I: Interpretations of the

Trusted Computer System Evaluation Criteria

Highlighting (ALL CAPS) is used in Part I to indicate criteria not contained in a lower class or changes and additions to already defined criteria. Where there is no highlighting, requirements have been carried over from lower classes without addition or modification.

1.0 DIVISION D: MINIMAL PROTECTION

This division contains only one class. It is reserved for those systems that have been evaluated but that fail to meet the requirements for a higher evaluation class.

2.0 DIVISION C: DISCRETIONARY PROTECTION

Classes in this division provide for discretionary (need-to-know) protection and, through the inclusion of audit capabilities, for accountability of subjects and the actions they initiate.

2.1 CLASS (C1): DISCRETIONARY SECURITY PROTECTION

THE NETWORK TRUSTED COMPUTING BASE (NTCB) OF A CLASS (C1) NETWORK SYSTEM NOMINALLY SATISFIES THE DISCRETIONARY SECURITY REQUIREMENTS BY PROVIDING SEPARATION OF USERS AND DATA. IT INCORPORATES SOME FORM OF CREDIBLE CONTROLS CAPABLE OF ENFORCING ACCESS LIMITATIONS ON AN INDIVIDUAL BASIS, I.E., OSTENSIBLY SUITABLE FOR ALLOWING USERS TO BE ABLE TO PROTECT PRIVATE OR PROJECT INFORMATION AND TO KEEP OTHER USERS FROM ACCIDENTALLY READING OR

DESTROYING THEIR DATA. THE CLASS (C1) ENVIRONMENT IS EXPECTED TO BE ONE OF COOPERATING USERS PROCESSING DATA AT THE SAME LEVEL(S) OF SENSITIVITY. THE FOLLOWING ARE MINIMAL REQUIREMENTS FOR SYSTEMS ASSIGNED A CLASS (C1) RATING.

2.1.1 Security Policy

+ Statement from DoD 5200.28-STD

Implied from the Introduction to the TCSEC.

+ Interpretation

THE NETWORK SPONSOR SHALL DESCRIBE THE OVERALL NETWORK SECURITY POLICY ENFORCED BY THE NTCB. AT A MINIMUM, THIS POLICY SHALL INCLUDE THE DISCRETIONARY REQUIREMENTS APPLICABLE TO THIS CLASS. THE POLICY MAY REQUIRE DATA SECRECY, OR DATA INTEGRITY, OR BOTH. THE POLICY SHALL INCLUDE A DISCRETIONARY POLICY FOR PROTECTING THE INFORMATION BEING PROCESSED BASED ON THE AUTHORIZATIONS OF USERS OR GROUPS OF USERS. THIS ACCESS CONTROL POLICY STATEMENT SHALL DESCRIBE THE REQUIREMENTS ON THE NETWORK TO PREVENT OR DETECT "READING OR DESTROYING" SENSITIVE INFORMATION BY UNAUTHORIZED USERS OR ERRORS. UNAUTHORIZED USERS INCLUDE BOTH THOSE THAT ARE NOT AUTHORIZED TO USE THE NETWORK AT ALL (E.G., A USER ATTEMPTING TO USE A PASSIVE OR ACTIVE WIRE TAP) OR A LEGITIMATE USER OF THE NETWORK WHO IS NOT AUTHORIZED TO ACCESS A SPECIFIC PIECE OF INFORMATION BEING PROTECTED.

NOTE THAT "USERS" DOES NOT INCLUDE "OPERATORS," "SYSTEM PROGRAMMERS," "TECHNICAL CONTROL OFFICERS," "SYSTEM SECURITY OFFICERS," AND OTHER SYSTEM SUPPORT PERSONNEL. THEY ARE DISTINCT FROM USERS AND ARE SUBJECT TO THE TRUSTED FACILITY MANUAL AND THE SYSTEM ARCHITECTURE REQUIREMENTS. SUCH INDIVIDUALS MAY CHANGE THE SYSTEM PARAMETERS OF THE NETWORK SYSTEM, FOR EXAMPLE, BY DEFINING MEMBERSHIP OF A GROUP. THESE INDIVIDUALS MAY ALSO HAVE THE SEPARATE ROLE OF USERS.

SECRECY POLICY: THE NETWORK SPONSOR SHALL DEFINE THE FORM OF THE DISCRETIONARY SECRECY POLICY THAT IS ENFORCED IN THE NETWORK TO PREVENT UNAUTHORIZED USERS FROM READING THE SENSITIVE INFORMATION ENTRUSTED TO THE NETWORK.

DATA INTEGRITY POLICY: THE NETWORK SPONSOR SHALL DEFINE THE DISCRETIONARY INTEGRITY POLICY TO PREVENT UNAUTHORIZED USERS FROM MODIFYING, VIZ., WRITING, SENSITIVE INFORMATION. THE DEFINITION OF DATA

INTEGRITY PRESENTED BY THE NETWORK SPONSOR REFERS TO THE REQUIREMENT THAT THE INFORMATION HAS NOT BEEN SUBJECTED TO UNAUTHORIZED MODIFICATION IN THE NETWORK.

+ Rationale

THE WORD "SPONSOR" IS USED IN PLACE OF ALTERNATIVES (SUCH AS "VENDOR," "ARCHITECT," "MANUFACTURER," AND "DEVELOPER") BECAUSE THE ALTERNATIVES INDICATE PEOPLE WHO MAY NOT BE AVAILABLE, INVOLVED, OR RELEVANT AT THE TIME THAT A NETWORK SYSTEM IS PROPOSED FOR EVALUATION.

A TRUSTED NETWORK IS ABLE TO CONTROL BOTH THE READING AND WRITING OF SHARED SENSITIVE INFORMATION. CONTROL OF WRITING IS USED TO PROTECT AGAINST DESTRUCTION OF INFORMATION. A NETWORK NORMALLY IS EXPECTED TO HAVE POLICY REQUIREMENTS TO PROTECT BOTH THE SECRECY AND INTEGRITY OF THE INFORMATION ENTRUSTED TO IT. IN A NETWORK THE INTEGRITY IS FREQUENTLY AS IMPORTANT OR MORE IMPORTANT THAN THE SECRECY REQUIREMENTS. THEREFORE THE SECRECY AND/OR INTEGRITY POLICY TO BE ENFORCED BY THE NETWORK MUST BE STATED FOR EACH NETWORK REGARDLESS OF ITS EVALUATION CLASS. THE ASSURANCE THAT THE POLICY IS FAITHFULLY ENFORCED IS REFLECTED IN THE EVALUATION CLASS OF THE NETWORK.

THIS CONTROL OVER MODIFICATION IS TYPICALLY USED TO PROTECT INFORMATION SO THAT IT MAY BE RELIED UPON AND TO CONTROL THE POTENTIAL HARM THAT WOULD RESULT IF THE INFORMATION WERE CORRUPTED. THE OVERALL NETWORK POLICY REQUIREMENTS FOR INTEGRITY INCLUDES THE PROTECTION FOR DATA BOTH WHILE BEING PROCESSED IN A COMPONENT AND WHILE BEING TRANSMITTED IN THE NETWORK. THE ACCESS CONTROL POLICY ENFORCED BY THE NTCB RELATES TO THE ACCESS OF SUBJECTS TO OBJECTS WITHIN EACH COMPONENT. COMMUNICATIONS INTEGRITY ADDRESSED WITHIN PART II RELATES TO INFORMATION WHILE BEING TRANSMITTED.

2.1.1.1 Discretionary Access Control

+ Statement from DoD 5200.28-STD

THE TCB SHALL DEFINE AND CONTROL ACCESS BETWEEN NAMED USERS AND NAMED OBJECTS (E.G., FILES AND PROGRAMS) IN THE ADP SYSTEM. THE ENFORCEMENT MECHANISM (E.G., SELF/GROUP/PUBLIC CONTROLS, ACCESS CONTROL LISTS) SHALL ALLOW USERS TO SPECIFY AND CONTROL SHARING OF THOSE OBJECTS BY NAMED INDIVIDUALS OR DEFINED GROUPS OF INDIVIDUALS, OR BOTH.

+ Interpretation

THE DISCRETIONARY ACCESS CONTROL (DAC) MECHANISM(S) MAY BE DISTRIBUTED OVER THE PARTITIONED NTCB IN VARIOUS WAYS. SOME PART, ALL, OR NONE OF THE DAC MAY BE IMPLEMENTED IN A GIVEN COMPONENT OF THE NETWORK SYSTEM. IN PARTICULAR, COMPONENTS THAT SUPPORT ONLY INTERNAL SUBJECTS (I.E., THAT HAVE NO SUBJECTS ACTING AS DIRECT SURROGATES FOR USERS), SUCH AS A PUBLIC NETWORK PACKET SWITCH, MIGHT NOT IMPLEMENT THE DAC MECHANISM(S) DIRECTLY (E.G., THEY ARE UNLIKELY TO CONTAIN ACCESS CONTROL LISTS).

IDENTIFICATION OF USERS BY GROUPS MAY BE ACHIEVED IN VARIOUS WAYS IN THE NETWORKING ENVIRONMENT. FOR EXAMPLE, THE NETWORK IDENTIFIERS (E.G., INTERNET ADDRESSES) FOR VARIOUS COMPONENTS (E.G., HOSTS, GATEWAYS) CAN BE USED AS IDENTIFIERS OF GROUPS OF INDIVIDUAL USERS (E.G., "ALL USERS AT HOST A," "ALL USERS OF NETWORK Q") WITHOUT EXPLICIT IDENTIFICATION OF INDIVIDUAL USERS, NOR EVEN AN EXPLICIT NUMBER OF USERS IMPLIED), IF THIS IS CONSISTENT WITH THE NETWORK SECURITY POLICY.

FOR NETWORKS, INDIVIDUAL HOSTS WILL IMPOSE NEED-TO-KNOW CONTROLS OVER THEIR USERS - MUCH LIKE (IN FACT, PROBABLY THE SAME) CONTROLS USED WHEN THERE IS NO NETWORK CONNECTION.

WHEN GROUP IDENTIFIERS ARE ACCEPTABLE FOR ACCESS CONTROL, THE IDENTIFIER OF SOME OTHER HOST MAY BE EMPLOYED, TO ELIMINATE THE MAINTENANCE THAT WOULD BE REQUIRED IF INDIVIDUAL IDENTIFICATION OF REMOTE USERS WAS EMPLOYED.

THE DAC MECHANISM OF A NTCB PARTITION MAY BE IMPLEMENTED AT THE INTERFACE OF THE REFERENCE MONITOR OR MAY BE DISTRIBUTED IN SUBJECTS THAT ARE PART OF THE NTCB IN THE SAME OR DIFFERENT COMPONENT. THE REFERENCE MONITOR MANAGES ALL THE PHYSICAL RESOURCES OF THE SYSTEM AND FROM THEM CREATES THE ABSTRACTION OF SUBJECTS AND OBJECTS THAT IT CONTROLS. SOME OF THESE SUBJECTS AND OBJECTS MAY BE USED TO IMPLEMENT A PART OF THE NTCB.

WHEN INTEGRITY IS INCLUDED AS PART OF THE NETWORK DISCRETIONARY SECURITY POLICY, THE ABOVE INTERPRETATIONS SHALL BE SPECIFICALLY APPLIED TO THE CONTROLS OVER MODIFICATION, VIZ, THE WRITE MODE OF ACCESS, WITHIN EACH COMPONENT BASED ON IDENTIFIED USERS OR GROUPS OF USERS.

+ Rationale

IN THIS CLASS, THE SUPPORTING ELEMENTS OF THE OVERALL DAC MECHANISM ARE TREATED EXACTLY AS UNTRUSTED SUBJECTS ARE TREATED WITH RESPECT TO DAC IN AN ADP SYSTEM, WITH THE SAME RESULT AS NOTED IN THE INTERPRETATION. STRENGTHENING OF THE DAC MECHANISM IN THE NETWORK ENVIRONMENT IS PROVIDED IN

CLASS (C2) (SEE THE DISCRETIONARY ACCESS CONTROL SECTION).

A TYPICAL SITUATION FOR DAC IS THAT A SURROGATE PROCESS FOR A REMOTE USER WILL BE CREATED IN SOME HOST FOR ACCESS TO OBJECTS UNDER THE CONTROL OF THE NTCB PARTITION WITHIN THAT HOST. THE INTERPRETATION REQUIRES THAT A USER IDENTIFIER BE ASSIGNED AND MAINTAINED FOR EACH SUCH PROCESS BY THE NTCB, SO THAT ACCESS BY A SURROGATE PROCESS IS SUBJECT TO ESSENTIALLY THE SAME DISCRETIONARY CONTROLS AS ACCESS BY A PROCESS ACTING ON BEHALF OF A LOCAL USER WOULD BE. HOWEVER, WITHIN THIS INTERPRETATION A RANGE OF POSSIBLE INTERPRETATIONS OF THE ASSIGNED USER IDENTIFICATION IS PERMITTED.

THE MOST OBVIOUS SITUATION WOULD EXIST IF A GLOBAL DATABASE OF NETWORK USERS WERE TO BE MADE PERMANENTLY AVAILABLE ON DEMAND TO EVERY HOST, (I.E., A NAME SERVER EXISTED) SO THAT ALL USER IDENTIFICATIONS WERE GLOBALLY MEANINGFUL.

IT IS ALSO ACCEPTABLE, HOWEVER, FOR SOME NTCB PARTITIONS TO MAINTAIN A DATABASE OF LOCALLY-REGISTERED USERS FOR ITS OWN USE. IN SUCH A CASE, ONE COULD CHOOSE TO INHIBIT THE CREATION OF SURROGATE PROCESSES FOR LOCALLY UNREGISTERED USERS, OR (IF PERMITTED BY THE LOCAL POLICY) ALTERNATIVELY, TO PERMIT THE CREATION OF SURROGATE PROCESSES WITH PRESELECTED USER AND GROUP IDENTIFIERS WHICH, IN EFFECT, IDENTIFY THE PROCESS AS EXECUTING ON BEHALF OF A MEMBER OF A GROUP OF USERS ON A PARTICULAR REMOTE HOST. THE INTENT OF THE WORDS CONCERNING AUDIT IN THE INTERPRETATION IS TO PROVIDE A MINIMALLY ACCEPTABLE DEGREE OF AUDITABILITY FOR CASES SUCH AS THE LAST DESCRIBED. WHAT IS REQUIRED IS THAT THERE BE A CAPABILITY, USING THE AUDIT FACILITIES PROVIDED BY THE NETWORK NTCB PARTITIONS INVOLVED, TO DETERMINE WHO WAS LOGGED IN AT THE ACTUAL HOST OF THE GROUP OF REMOTE USERS AT THE TIME THE SURROGATE PROCESSING OCCURED.

ASSOCIATING THE PROPER USER ID WITH A SURROGATE PROCESS IS THE JOB OF IDENTIFICATION AND AUTHENTICATION. THIS MEANS THAT DAC IS APPLIED LOCALLY, WITH RESPECT TO THE USER ID OF THE SURROGATE PROCESS. THE TRANSMISSION OF THE DATA BACK ACROSS THE NETWORK TO THE USER'S HOST, AND THE CREATION OF A COPY OF THE DATA THERE, IS NOT THE BUSINESS OF DAC.

COMPONENTS THAT SUPPORT ONLY INTERNAL SUBJECTS IMPACT THE IMPLEMENTATION OF THE DAC BY PROVIDING SERVICES BY WHICH INFORMATION (E.G., A USER-ID) IS MADE AVAILABLE TO A COMPONENT THAT MAKES A DAC DECISION. AN EXAMPLE OF THE LATTER WOULD BE THE CASE THAT A USER AT HOST A ATTEMPTS TO ACCESS A FILE AT HOST B. THE DAC DECISION MIGHT BE (AND USUALLY WOULD BE) MADE AT HOST B ON THE BASIS OF A USER-ID TRANSMITTED FROM HOST A TO HOST B.

UNIQUE USER IDENTIFICATION MAY BE ACHIEVED BY A VARIETY OF MECHANISMS, INCLUDING (A) A REQUIREMENT FOR UNIQUE IDENTIFICATION AND AUTHENTICATION ON THE HOST WHERE ACCESS TAKES

PLACE; (B) RECOGNITION OF FULLY QUALIFIED NETWORK ADDRESSES AUTHENTICATED BY ANOTHER HOST AND FORWARDED TO THE HOST WHERE ACCESS TAKES PLACE; OR (C) ADMINISTRATIVE SUPPORT OF A NETWORK-WIDE UNIQUE PERSONNEL IDENTIFIER THAT COULD BE AUTHENTICATED AND FORWARDED BY ANOTHER HOST AS IN (B) ABOVE, OR COULD BE AUTHENTICATED AND FORWARDED BY A DEDICATED NETWORK IDENTIFICATION AND AUTHENTICATION SERVER. THE PROTOCOLS WHICH IMPLEMENT (B) OR (C) ARE SUBJECT TO THE SYSTEM ARCHITECTURE REQUIREMENTS.

NETWORK SUPPORT FOR DAC MIGHT BE HANDLED IN OTHER WAYS THAN THAT DESCRIBED AS "TYPICAL" ABOVE. IN PARTICULAR, SOME FORM OF CENTRALIZED ACCESS CONTROL IS OFTEN PROPOSED. AN ACCESS CONTROL CENTER MAY MAKE ALL DECISIONS FOR DAC, OR IT MAY SHARE THE BURDEN WITH THE HOSTS BY CONTROLLING HOST-TO-HOST CONNECTIONS, AND LEAVING THE HOSTS TO DECIDE ON ACCESS TO THEIR OBJECTS BY USERS AT A LIMITED SET OF REMOTE HOSTS. IN THIS CASE THE ACCESS CONTROL CENTER PROVIDES THE LINKAGE BETWEEN THE CONNECTION ORIENTED ABSTRACTION (AS DISCUSSED IN THE INTRODUCTION) AND THE OVERALL NETWORK SECURITY POLICY FOR DAC. IN ALL CASES THE ENFORCEMENT OF THE DECISION MUST BE PROVIDED BY THE HOST WHERE THE OBJECT RESIDES.

2.1.2 Accountability

2.1.2.1 Identification and Authentication

+ Statement from DoD 5200.28-STD

THE TCB SHALL REQUIRE USERS TO IDENTIFY THEMSELVES TO IT BEFORE BEGINNING TO PERFORM ANY OTHER ACTIONS THAT THE TCB IS EXPECTED TO MEDIATE. FURTHERMORE, THE TCB SHALL USE A PROTECTED MECHANISM (E.G., PASSWORDS) TO AUTHENTICATE THE USER'S IDENTITY. THE TCB SHALL PROTECT AUTHENTICATION DATA SO THAT IT CANNOT BE ACCESSED BY ANY UNAUTHORIZED USER.

+ Interpretation

THE REQUIREMENT FOR IDENTIFICATION AND AUTHENTICATION OF USERS IS THE SAME FOR A NETWORK SYSTEM AS FOR AN ADP SYSTEM. THE IDENTIFICATION AND AUTHENTICATION MAY BE DONE BY THE COMPONENT TO WHICH THE USER IS DIRECTLY CONNECTED OR SOME OTHER COMPONENT, SUCH AS AN IDENTIFICATION AND AUTHENTICATION SERVER. AVAILABLE TECHNIQUES, SUCH AS THOSE DESCRIBED IN THE PASSWORD GUIDELINE=, ARE GENERALLY ALSO APPLICABLE IN THE NETWORK CONTEXT. HOWEVER, IN CASES WHERE THE NTCB IS EXPECTED TO MEDIATE ACTIONS OF A HOST (OR OTHER NETWORK COMPONENT) THAT IS ACTING ON BEHALF OF A USER OR GROUP OF USERS, THE NTCB MAY EMPLOY IDENTIFICATION AND AUTHENTICATION OF THE HOST (OR OTHER COMPONENT) IN LIEU OF IDENTIFICATION AND AUTHENTICATION OF AN INDIVIDUAL USER.

AUTHENTICATION INFORMATION, INCLUDING THE IDENTITY OF A USER (ONCE AUTHENTICATED) MAY BE PASSED FROM ONE COMPONENT TO ANOTHER WITHOUT REAUTHENTICATION, SO LONG AS THE NTCB PROTECTS (E.G., BY ENCRYPTION) THE INFORMATION FROM UNAUTHORIZED DISCLOSURE AND MODIFICATION. THIS PROTECTION SHALL PROVIDE AT LEAST A SIMILAR LEVEL OF ASSURANCE (OR STRENGTH OF MECHANISM) AS PERTAINS TO THE PROTECTION OF THE AUTHENTICATION MECHANISM AND AUTHENTICATION DATA.

+ Rationale

THE NEED FOR ACCOUNTABILITY IS NOT CHANGED IN THE CONTEXT OF A NETWORK SYSTEM. THE FACT THAT THE NTCB IS PARTITIONED OVER A SET OF COMPONENTS NEITHER REDUCES THE NEED NOR IMPOSES NEW REQUIREMENTS. THAT IS, INDIVIDUAL ACCOUNTABILITY IS STILL THE OBJECTIVE. HOWEVER, IN THE CONTEXT OF A NETWORK SYSTEM AT THE (C1) LEVEL (WHEREIN EXPLICIT INDIVIDUAL USER ACCOUNTABILITY IS NOT REQUIRED), "INDIVIDUAL ACCOUNTABILITY" CAN BE SATISFIED BY IDENTIFICATION OF A HOST (OR OTHER COMPONENT). IN ADDITION, THERE IS NO NEED IN A DISTRIBUTED PROCESSING SYSTEM LIKE A NETWORK TO REAUTHENTICATE A USER AT EACH POINT IN THE NETWORK WHERE A PROJECTION OF A USER (VIA THE SUBJECT OPERATING ON BEHALF OF THE USER) INTO ANOTHER REMOTE SUBJECT TAKES PLACE.

THE PASSING OF IDENTIFIERS AND/OR AUTHENTICATION INFORMATION FROM ONE COMPONENT TO ANOTHER IS USUALLY DONE IN SUPPORT TO THE IMPLEMENTATION OF THE DISCRETIONARY ACCESS CONTROL (DAC). THIS SUPPORT RELATES DIRECTLY TO THE DAC REGARDING ACCESS BY A USER TO A STORAGE OBJECT IN A DIFFERENT NTCB PARTITION THAN THE ONE WHERE THE USER WAS AUTHENTICATED. EMPLOYING A FORWARDED IDENTIFICATION IMPLIES ADDITIONAL RELIANCE ON THE SOURCE AND COMPONENTS ALONG THE PATH.

2.1.3 Assurance

2.1.3.1 Operational Assurance

2.1.3.1.1 System Architecture

+ Statement from DoD 5200.28-STD

THE TCB SHALL MAINTAIN A DOMAIN FOR ITS OWN EXECUTION THAT PROTECTS IT FROM EXTERNAL INTERFERENCE OR TAMPERING (E.G., BY MODIFICATION OF ITS CODE OR DATA STRUCTURES). RESOURCES CONTROLLED BY THE TCB MAY BE A DEFINED SUBSET OF THE SUBJECTS AND OBJECTS IN THE ADP SYSTEM.

+ Interpretation

THE SYSTEM ARCHITECTURE CRITERION MUST BE MET INDIVIDUALLY BY ALL NTCB PARTITIONS. IMPLEMENTATION OF THE REQUIREMENT THAT THE NTCB MAINTAIN A DOMAIN FOR ITS OWN EXECUTION IS ACHIEVED BY HAVING EACH NTCB PARTITION MAINTAIN A DOMAIN FOR ITS OWN EXECUTION.

THE SUBSET OF NETWORK RESOURCES OVER WHICH THE NTCB HAS CONTROL ARE THE UNION OF THE SETS OF RESOURCES OVER WHICH THE NTCB PARTITIONS HAVE CONTROL. CODE AND DATA STRUCTURES BELONGING TO THE NTCB, TRANSFERRED AMONG NTCB SUBJECTS (I.E., SUBJECTS OUTSIDE THE REFERENCE MONITOR BUT INSIDE THE NTCB) BELONGING TO DIFFERENT NTCB PARTITIONS, MUST BE PROTECTED AGAINST EXTERNAL INTERFERENCE OR TAMPERING. FOR EXAMPLE, A CRYPTOGRAPHIC CHECKSUM OR PHYSICAL MEANS MAY BE EMPLOYED TO PROTECT USER AUTHENTICATION DATA EXCHANGED BETWEEN NTCB PARTITIONS.

+ Rationale

THE REQUIREMENT FOR THE PROTECTION OF COMMUNICATIONS BETWEEN NTCB PARTITIONS IS SPECIFICALLY DIRECTED TO SUBJECTS THAT ARE PART OF THE NTCB PARTITIONS. ANY REQUIREMENTS FOR SUCH PROTECTION FOR THE SUBJECTS THAT ARE OUTSIDE THE NTCB PARTITIONS ARE ADDRESSED IN RESPONSE TO THE INTEGRITY REQUIREMENTS OF THE SECURITY POLICY.

2.1.3.1.2 System Integrity

+ Statement from DoD 5200.28-STD

HARDWARE AND/OR SOFTWARE FEATURES SHALL BE PROVIDED THAT CAN BE USED TO PERIODICALLY VALIDATE THE CORRECT OPERATION OF THE ON-SITE HARDWARE AND FIRMWARE ELEMENTS OF THE TCB.

+ Interpretation

IMPLEMENTATION OF THE REQUIREMENT IS PARTLY ACHIEVED BY HAVING HARDWARE AND/OR SOFTWARE FEATURES THAT CAN BE USED TO PERIODICALLY VALIDATE THE CORRECT OPERATION OF THE HARDWARE AND FIRMWARE ELEMENTS OF EACH COMPONENT'S NTCB PARTITION. FEATURES SHALL ALSO BE PROVIDED TO VALIDATE THE IDENTITY AND CORRECT OPERATION OF A COMPONENT PRIOR TO ITS INCORPORATION IN THE NETWORK SYSTEM AND THROUGHOUT SYSTEM OPERATION. FOR EXAMPLE, A PROTOCOL COULD BE DESIGNED THAT ENABLES THE COMPONENTS OF THE PARTITIONED NTCB TO EXCHANGE MESSAGES PERIODICALLY AND VALIDATE EACH OTHER'S CORRECT RESPONSE. THE PROTOCOL SHALL BE ABLE TO DETERMINE THE REMOTE ENTITY'S ABILITY TO RESPOND. NTCB PARTITIONS SHALL PROVIDE THE CAPABILITY TO REPORT TO NETWORK ADMINISTRATIVE PERSONNEL THE FAILURES DETECTED IN OTHER NTCB PARTITIONS.

INTERCOMPONENT PROTOCOLS IMPLEMENTED WITHIN A NTCB SHALL BE DESIGNED IN SUCH A WAY AS TO PROVIDE CORRECT OPERATION IN THE CASE OF FAILURES OF NETWORK COMMUNICATIONS OR INDIVIDUAL COMPONENTS. THE ALLOCATION OF DISCRETIONARY ACCESS CONTROL POLICY IN A NETWORK MAY REQUIRE COMMUNICATION BETWEEN TRUSTED SUBJECTS THAT ARE PART OF THE NTCB PARTITIONS IN DIFFERENT COMPONENTS. THIS COMMUNICATION IS NORMALLY IMPLEMENTED WITH A PROTOCOL BETWEEN THE SUBJECTS AS PEER ENTITIES. INCORRECT ACCESS WITHIN A COMPONENT SHALL NOT RESULT FROM FAILURE OF AN NTCB PARTITION TO COMMUNICATE WITH OTHER COMPONENTS.

+ Rationale

THE FIRST PARAGRAPH OF THE INTERPRETATION IS A STRAIGHTFORWARD EXTENSION OF THE REQUIREMENT INTO THE CONTEXT OF A NETWORK SYSTEM AND PARTITIONED NTCB AS DEFINED FOR THESE NETWORK CRITERIA.

NTCB PROTOCOLS SHOULD BE ROBUST ENOUGH SO THAT THEY PERMIT THE SYSTEM TO OPERATE CORRECTLY IN THE CASE OF LOCALIZED FAILURE. THE PURPOSE OF THIS PROTECTION IS TO PRESERVE THE INTEGRITY OF THE NTCB ITSELF. IT IS NOT UNUSUAL FOR ONE OR MORE COMPONENTS IN A NETWORK TO BE INOPERATIVE AT ANY TIME, SO IT IS IMPORTANT TO MINIMIZE THE EFFECTS OF SUCH FAILURES ON THE REST OF THE NETWORK. ADDITIONAL INTEGRITY AND DENIAL OF SERVICE ISSUES ARE ADDRESSED IN PART II.

2.1.3.2 Life-Cycle Assurance

2.1.3.2.1 Security Testing

+ Statement from DoD 5200.28-STD

THE SECURITY MECHANISMS OF THE ADP SYSTEM SHALL BE TESTED AND FOUND TO WORK AS CLAIMED IN THE SYSTEM DOCUMENTATION. TESTING SHALL BE DONE TO ASSURE THAT THERE ARE NO OBVIOUS WAYS FOR AN UNAUTHORIZED USER TO BYPASS OR OTHERWISE DEFEAT THE SECURITY PROTECTION MECHANISMS OF THE TCB. (SEE THE SECURITY TESTING GUIDELINES.)

+ Interpretation

TESTING OF A COMPONENT WILL REQUIRE A TESTBED THAT EXERCISES THE INTERFACES AND PROTOCOLS OF THE COMPONENT. THE TESTING OF A SECURITY MECHANISM OF THE NETWORK SYSTEM FOR MEETING THIS CRITERION SHALL BE AN INTEGRATED TESTING PROCEDURE INVOLVING ALL COMPONENTS CONTAINING AN NTCB PARTITION THAT IMPLEMENT THE GIVEN MECHANISM. THIS INTEGRATED TESTING IS ADDITIONAL TO ANY INDIVIDUAL COMPONENT TESTS INVOLVED IN THE EVALUATION OF THE NETWORK SYSTEM. THE SPON-

SOR SHOULD IDENTIFY THE ALLOWABLE SET OF CONFIGURATIONS INCLUDING THE SIZES OF THE NETWORKS. ANALYSIS OR TESTING PROCEDURES AND TOOLS SHALL BE AVAILABLE TO TEST THE LIMITS OF THESE CONFIGURATIONS. A CHANGE IN CONFIGURATION WITHIN THE ALLOWABLE SET OF CONFIGURATIONS DOES NOT REQUIRE RETESTING.

+ Rationale

TESTING IS THE PRIMARY METHOD AVAILABLE IN THIS EVALUATION DIVISION TO GAIN ANY ASSURANCE THAT THE SECURITY MECHANISMS PERFORM THEIR INTENDED FUNCTION.

2.1.4 Documentation

2.1.4.1 Security Features User's Guide

+ Statement from DoD 5200.28-STD

A SINGLE SUMMARY, CHAPTER, OR MANUAL IN USER DOCUMENTATION SHALL DESCRIBE THE PROTECTION MECHANISMS PROVIDED BY THE TCB, INTERPRETATIONS ON THEIR USE, AND HOW THEY INTERACT WITH ONE ANOTHER.

+ Interpretation

THIS USER DOCUMENTATION DESCRIBES USER VISIBLE PROTECTION MECHANISMS AT THE GLOBAL (NETWORK SYSTEM) LEVEL AND AT THE USER INTERFACE OF EACH COMPONENT, AND THE INTERACTION AMONG THESE.

+ Rationale

THE INTERPRETATION IS AN EXTENSION OF THE REQUIREMENT INTO THE CONTEXT OF A NETWORK SYSTEM AS DEFINED FOR THESE NETWORK CRITERIA. DOCUMENTATION OF PROTECTION MECHANISMS PROVIDED BY INDIVIDUAL COMPONENTS IS REQUIRED BY THE CRITERIA FOR TRUSTED COMPUTER SYSTEMS THAT ARE APPLIED AS APPROPRIATE FOR THE INDIVIDUAL COMPONENTS.

2.1.4.2 Trusted Facility Manual

+ Statement from DoD 5200.28-STD

A MANUAL ADDRESSED TO THE ADP SYSTEM ADMINISTRATOR SHALL PRESENT CAUTIONS ABOUT FUNCTIONS AND PRIVILEGES THAT SHOULD BE CONTROLLED WHEN RUNNING A SECURE FACILITY.

+ Interpretation

THIS MANUAL SHALL CONTAIN SPECIFICATIONS AND PROCEDURES TO ASSIST THE SYSTEM ADMINISTRATOR(S) MAINTAIN COGNIZANCE OF

THE NETWORK CONFIGURATION. THESE SPECIFICATIONS AND PROCEDURES SHALL ADDRESS THE FOLLOWING:

- 1.THE HARDWARE CONFIGURATION OF THE NETWORK ITSELF;
- 2.THE IMPLICATIONS OF ATTACHING NEW COMPONENTS TO THE NETWORK;
- 3.THE CASE WHERE CERTAIN COMPONENTS MAY PERIODICALLY LEAVE THE NETWORK (E.G., BY CRASHING, OR BY BEING DISCONNECTED) AND THEN REJOIN;
- 4.NETWORK CONFIGURATION ASPECTS THAT CAN IMPACT THE SECURITY OF THE NETWORK SYSTEM; (FOR EXAMPLE, THE MANUAL SHOULD DESCRIBE FOR THE NETWORK SYSTEM ADMINISTRATOR THE INTERCONNECTIONS AMONG COMPONENTS THAT ARE CONSISTENT WITH THE OVERALL NETWORK SYSTEM ARCHITECTURE.)
- 5.LOADING OR MODIFYING NTCB SOFTWARE OR FIRMWARE (E.G., DOWN-LINE LOADING).

THE PHYSICAL AND ADMINISTRATIVE ENVIRONMENTAL CONTROLS SHALL BE SPECIFIED. ANY ASSUMPTIONS ABOUT SECURITY OF A GIVEN NETWORK SHOULD BE CLEARLY STATED (E.G., THE FACT THAT ALL COMMUNICATIONS LINKS MUST BE PHYSICALLY PROTECTED TO A CERTAIN LEVEL).

+ Rationale

THERE MAY BE MULTIPLE SYSTEM ADMINISTRATORS WITH DIVERSE RESPONSIBILITIES. THE TECHNICAL SECURITY MEASURES DESCRIBED BY THESE CRITERIA MUST BE USED IN CONJUNCTION WITH OTHER FORMS OF SECURITY IN ORDER TO ACHIEVE SECURITY OF THE NETWORK. ADDITIONAL FORMS INCLUDE ADMINISTRATIVE SECURITY, PHYSICAL SECURITY, EMANATIONS SECURITY, ETC.

EXTENSION OF THIS CRITERION TO COVER CONFIGURATION ASPECTS OF THE NETWORK IS NEEDED BECAUSE, FOR EXAMPLE, PROPER INTERCONNECTION OF COMPONENTS IS TYPICALLY ESSENTIAL TO ACHIEVE A CORRECT REALIZATION OF THE NETWORK ARCHITECTURE.

CRYPTOGRAPHY IS ONE COMMON MECHANISM EMPLOYED TO PROTECT COMMUNICATION CIRCUITS. ENCRYPTION TRANSFORMS THE REPRESENTATION OF INFORMATION SO THAT IT IS UNINTELLIGIBLE TO UNAUTHORIZED SUBJECTS. REFLECTING THIS TRANSFORMATION, THE SENSITIVITY OF THE CIPHERTEXT IS GENERALLY LOWER THAN THE CLARTEXT. IF ENCRYPTION METHODOLOGIES ARE EMPLOYED, THEY SHALL BE APPROVED BY THE NATIONAL SECURITY AGENCY (NSA).

THE ENCRYPTION ALGORITHM AND ITS IMPLEMENTATION ARE OUTSIDE THE SCOPE OF THESE INTERPRETATIONS. THIS ALGORITHM AND IMPLEMENTATION MAY BE IMPLEMENTED IN A SEPARATE DEVICE OR MAY BE A FUNCTION OF A SUBJECT IN A COMPONENT NOT DEDICATED TO ENCRYPTION. WITHOUT PREJUDICE, EITHER IMPLEMENTATION PACKAGING IS REFERRED TO AS AN ENCRYPTION MECHANISM HEREIN.

2.1.4.3 Test Documentation

+ Statement from DoD 5200.28-STD

THE SYSTEM DEVELOPER SHALL PROVIDE TO THE EVALUATORS A DOCUMENT THAT DESCRIBES THE TEST PLAN, TEST PROCEDURES THAT SHOW HOW THE SECURITY MECHANISMS WERE TESTED, AND RESULTS OF THE SECURITY MECHANISMS' FUNCTIONAL TESTING.

+ Interpretation

THE "SYSTEM DEVELOPER" IS INTERPRETED AS "THE NETWORK SYSTEM SPONSOR". THE DESCRIPTION OF THE TEST PLAN SHOULD ESTABLISH THE CONTEXT IN WHICH THE TESTING WAS OR SHOULD BE CONDUCTED. THE DESCRIPTION SHOULD IDENTIFY ANY ADDITIONAL TEST COMPONENTS THAT ARE NOT PART OF THE SYSTEM BEING EVALUATED. THIS INCLUDES A DESCRIPTION OF THE TEST-RELEVANT FUNCTIONS OF SUCH TEST COMPONENTS AND A DESCRIPTION OF THE INTERFACING OF THOSE TEST COMPONENTS TO THE SYSTEM BEING EVALUATED. THE DESCRIPTION OF THE TEST PLAN SHOULD ALSO DEMONSTRATE THAT THE TESTS ADEQUATELY COVER THE NETWORK SECURITY POLICY. THE TESTS SHOULD INCLUDE THE FEATURES DESCRIBED IN THE SYSTEM ARCHITECTURE AND THE SYSTEM INTEGRITY SECTIONS. THE TESTS SHOULD ALSO INCLUDE NETWORK CONFIGURATION AND SIZING.

+ Rationale

THE ENTITY BEING EVALUATED MAY BE A NETWORKING SUBSYSTEM (SEE APPENDIX A) TO WHICH OTHER COMPONENTS MUST BE ADDED TO MAKE A COMPLETE NETWORK SYSTEM. IN THAT CASE, THIS INTERPRETATION IS EXTENDED TO INCLUDE CONTEXTUAL DEFINITION BECAUSE, AT EVALUATION TIME, IT IS NOT POSSIBLE TO VALIDATE THE TEST PLANS WITHOUT THE DESCRIPTION OF THE CONTEXT FOR TESTING THE NETWORKING SUBSYSTEM.

2.1.4.4 Design Documentation

+ Statement from DoD 5200.28-STD

DOCUMENTATION SHALL BE AVAILABLE THAT PROVIDES A DESCRIPTION OF THE MANUFACTURER'S PHILOSOPHY OF PROTECTION AND AN EXPLANATION OF HOW THIS PHILOSOPHY IS TRANSLATED INTO THE TCB. IF THE TCB IS COMPOSED OF DISTINCT MODULES, THE INTERFACES BETWEEN THESE MODULES SHALL BE DESCRIBED.

+ Interpretation

EXPLANATION OF HOW THE SPONSOR'S PHILOSOPHY OF PROTECTION IS TRANSLATED INTO THE NTCB SHALL INCLUDE A DESCRIPTION OF HOW THE NTCB IS PARTITIONED. THE SECURITY POLICY ALSO SHALL BE STATED. THE DESCRIPTION OF THE INTERFACES BETWEEN THE NTCB MODULES SHALL INCLUDE THE INTERFACE(S) BETWEEN NTCB PARTITIONS AND MODULES WITHIN THE PARTITIONS IF THE MODULES EXIST. THE SPONSOR SHALL DESCRIBE THE SECURITY ARCHITECTURE AND DESIGN, INCLUDING THE ALLOCATION OF SECURITY REQUIREMENTS AMONG COMPONENTS. APPENDIX A ADDRESSES COMPONENT EVALUATION ISSUES.

+ Rationale

THE INTERPRETATION IS A STRAIGHTFORWARD EXTENSION OF THE REQUIREMENT INTO THE CONTEXT OF A NETWORK SYSTEM AS DEFINED FOR THIS NETWORK INTERPRETATION. OTHER DOCUMENTATION, SUCH AS DESCRIPTION OF COMPONENTS AND DESCRIPTION OF OPERATING ENVIRONMENT(S) IN WHICH THE NETWORKING SUBSYSTEM OR NETWORK SYSTEM IS DESIGNED TO FUNCTION, IS REQUIRED ELSEWHERE, E.G., IN THE TRUSTED FACILITY MANUAL.

IN ORDER TO BE EVALUATED, A NETWORK MUST POSSESS A COHERENT NETWORK SECURITY ARCHITECTURE AND DESIGN. (INTERCONNECTION OF COMPONENTS THAT DO NOT ADHERE TO SUCH A SINGLE COHERENT NETWORK SECURITY ARCHITECTURE IS ADDRESSED IN THE INTERCONNECTION OF ACCREDITED AIS, APPENDIX C.) THE NETWORK SECURITY ARCHITECTURE MUST ADDRESS THE SECURITY-RELEVANT POLICIES, OBJECTIVES, AND PROTOCOLS. THE NETWORK SECURITY DESIGN SPECIFIES THE INTERFACES AND SERVICES THAT MUST BE INCORPORATED INTO THE NETWORK SO THAT IT CAN BE EVALUATED AS A TRUSTED ENTITY. THERE MAY BE MULTIPLE DESIGNS THAT CONFORM TO THE SAME ARCHITECTURE BUT ARE MORE OR LESS INCOMPATIBLE AND NON-INTEROPERABLE (EXCEPT THROUGH THE INTERCONNECTION RULES). SECURITY RELATED MECHANISMS REQUIRING COOPERATION AMONG COMPONENTS ARE SPECIFIED IN THE DESIGN IN TERMS OF THEIR VISIBLE INTERFACES; MECHANISMS HAVING NO VISIBLE INTERFACES ARE NOT SPECIFIED IN THIS DOCUMENT BUT ARE LEFT AS IMPLEMENTATION DECISIONS.

THE NETWORK SECURITY ARCHITECTURE AND DESIGN MUST BE AVAILABLE FROM THE NETWORK SPONSOR BEFORE EVALUATION OF THE NETWORK, OR ANY COMPONENT, CAN BE UNDERTAKEN. THE NETWORK SECURITY ARCHITECTURE AND DESIGN MUST BE SUFFICIENTLY COMPLETE, UNAMBIGUOUS, AND FREE FROM OBVIOUS FLAWS TO PERMIT THE CONSTRUCTION OR ASSEMBLY OF A TRUSTED NETWORK BASED ON THE STRUCTURE IT SPECIFIES.

WHEN A COMPONENT IS BEING DESIGNED OR PRESENTED FOR EVALUATION, OR WHEN A NETWORK ASSEMBLED FROM COMPONENTS IS ASSEMBLED OR PRESENTED FOR EVALUATION, THERE MUST BE A PRIORI EVIDENCE THAT THE NETWORK SECURITY ARCHITECTURE AND DESIGN ARE SATISFIED. THAT IS, THE COMPONENTS CAN BE ASSEMBLED INTO A NETWORK THAT CONFORMS IN EVERY WAY WITH THE NETWORK SECURITY ARCHITECTURE AND DESIGN TO PRODUCE A PHYSICAL REALIZATION THAT IS TRUSTED TO THE EXTENT THAT ITS EVALUATION INDICATES.

IN ORDER FOR A TRUSTED NETWORK TO BE CONSTRUCTED FROM COMPONENTS THAT CAN BE BUILT INDEPENDENTLY, THE NETWORK SECURITY ARCHITECTURE AND DESIGN MUST COMPLETELY AND UNAMBIGUOUSLY DEFINE THE SECURITY FUNCTIONALITY OF COMPONENTS AS WELL AS THE INTERFACES BETWEEN OR AMONG COMPONENTS. THE NETWORK SECURITY ARCHITECTURE AND DESIGN MUST BE EVALUATED TO DETERMINE THAT A NETWORK CONSTRUCTED TO ITS SPECIFICATIONS WILL IN FACT BE TRUSTED, THAT IS, IT WILL BE EVALUATABLE UNDER THESE INTERPRETATIONS.

2.2 CLASS (C2): CONTROLLED ACCESS PROTECTION

NETWORK SYSTEMS IN THIS CLASS ENFORCE A MORE FINELY GRAINED DISCRETIONARY ACCESS CONTROL THAN (C1) NETWORK SYSTEMS, MAKING USERS INDIVIDUALLY ACCOUNTABLE FOR THEIR ACTIONS THROUGH LOGIN PROCEDURES, AUDITING OF SECURITY-RELEVANT EVENTS, AND RESOURCE ISOLATION. THE FOLLOWING ARE MINIMAL REQUIREMENTS FOR SYSTEMS ASSIGNED A CLASS (C2) RATING.

2.2.1 Security Policy

+ Statement from DoD 5200.28-STD

Implied from the Introduction to the TCSEC.

+ Interpretation

The network sponsor shall describe the overall network security policy enforced by the NTCB. At a minimum, this policy shall include the discretionary requirements applicable to this class. The policy may require data secrecy, or data integrity, or both. The policy shall include a discretionary policy for protecting the information being pro-

cessed based on the authorizations of INDIVIDUALS, users, or groups of users. This access control policy statement shall describe the requirements on the network to prevent or detect "reading or destroying" sensitive information by unauthorized users or errors. Unauthorized users include both those that are not authorized to use the network at all (e.g., a user attempting to use a passive or active wire tap) or a legitimate user of the network who is not authorized to access a specific piece of information being protected.

Note that "users" does not include "operators," "system programmers," "technical control officers," "system security officers," and other system support personnel. They are distinct from users and are subject to the Trusted Facility Manual and the System Architecture requirements. Such individuals may change the system parameters of the network system, for example, by defining membership of a group. These individuals may also have the separate role of users.

SECURITY POLICY: The network sponsor shall define the form of the discretionary security policy that is enforced in the network to prevent unauthorized users from reading the sensitive information entrusted to the network.

DATA INTEGRITY POLICY: The network sponsor shall define the discretionary integrity policy to prevent unauthorized users from modifying, viz., writing, sensitive information. The definition of data integrity presented by the network sponsor refers to the requirement that the information has not been subjected to unauthorized modification in the network.

+ Rationale

The word "sponsor" is used in place of alternatives (such as "vendor," "architect," "manufacturer," and "developer") because the alternatives indicate people who may not be available, involved, or relevant at the time that a network system is proposed for evaluation.

A trusted network is able to control both the reading and writing of shared sensitive information. Control of writing is used to protect against destruction of information. A network normally is expected to have policy requirements to protect both the security and integrity of the information entrusted to it. In a network the integrity is frequently as important or more important than the security requirements. Therefore the security and/or integrity policy to be enforced by the network must be stated for each net-

work regardless of its evaluation class. The assurance that the policy is faithfully enforced is reflected in the evaluation class of the network.

This control over modification is typically used to protect information so that it may be relied upon and to control the potential harm that would result if the information were corrupted. The overall network policy requirements for integrity includes the protection for data both while being processed in a component and while being transmitted in the network. The access control policy enforced by the NTCB relates to the access of subjects to objects within each component. Communications integrity addressed within Part II relates to information while being transmitted.

2.2.1.1 Discretionary Access Control

+ Statement from DoD 5200.28-STD

The TCB shall define and control access between named users and named objects (e.g., files and programs) in the ADP system. The enforcement mechanism (e.g., self/group/public controls, access control lists) shall allow users to specify and control sharing of those objects by named individuals or defined groups OF INDIVIDUALS, or both, AND SHALL PROVIDE CONTROLS TO LIMIT PROPAGATION OF ACCESS RIGHTS. THE DISCRETIONARY ACCESS CONTROL MECHANISM SHALL, EITHER BY EXPLICIT USER ACTION OR BY DEFAULT, PROVIDE THAT OBJECTS ARE PROTECTED FROM UNAUTHORIZED ACCESS. THESE ACCESS CONTROLS SHALL BE CAPABLE OF INCLUDING OR EXCLUDING ACCESS TO THE GRANULARITY OF A SINGLE USER. ACCESS PERMISSION TO AN OBJECT BY USERS NOT ALREADY POSSESSING ACCESS PERMISSION SHALL ONLY BE ASSIGNED BY AUTHORIZED USERS.

+ Interpretation

The discretionary access control (DAC) mechanism(s) may be distributed over the partitioned NTCB in various ways. Some part, all, or none of the DAC may be implemented in a given component of the network system. In particular, components that support only internal subjects (i.e., that have no subjects acting as direct surrogates for users), such as a public network packet switch, might not implement the DAC mechanism(s) directly (e.g., they are unlikely to contain access control lists).

Identification of users by groups may be achieved in various ways in the networking environment. For example, the network identifiers (e.g., internet addresses) for various components (e.g., hosts, gateways) can be used as identifiers of groups of individual users (e.g., "all users at Host A," "all users of network Q") SO LONG AS THE INDIVIDUALS INVOLVED IN THE GROUP ARE IMPLIED BY THE GROUP IDENTIF-

IER. FOR EXAMPLE, HOST A MIGHT EMPLOY A PARTICULAR GROUP-ID, FOR WHICH IT MAINTAINS A LIST OF EXPLICIT USERS IN THAT GROUP, IN ITS NETWORK EXCHANGE WITH HOST B, WHICH ACCEPTS THE GROUP-ID UNDER THE CONDITIONS OF THIS INTERPRETATION.

For networks, individual hosts will impose need-to-know controls over their users ON THE BASIS OF NAMED INDIVIDUALS - much like (in fact, probably the same) controls used when there is no network connection.

When group identifiers are acceptable for access control, the identifier of some other host may be employed, to eliminate the maintenance that would be required if individual identification of remote users was employed. IN CLASS C2 AND HIGHER, HOWEVER, IT MUST BE POSSIBLE FROM THAT AUDIT RECORD TO IDENTIFY (IMMEDIATELY OR AT SOME LATER TIME) EXACTLY THE INDIVIDUALS REPRESENTED BY A GROUP IDENTIFIER AT THE TIME OF THE USE OF THAT IDENTIFIER. THERE IS ALLOWED TO BE AN UNCERTAINTY BECAUSE OF ELAPSED TIME BETWEEN CHANGES IN THE GROUP MEMBERSHIP AND THE ENFORCEMENT IN THE ACCESS CONTROL MECHANISMS.

The DAC mechanism of a NTCB partition may be implemented at the interface of the reference monitor or may be distributed in subjects that are part of the NTCB in the same or different component. The reference monitor manages all the physical resources of the system and from them creates the abstraction of subjects and objects that it controls. Some of these subjects and objects may be used to implement a part of the NTCB. WHEN THE DAC MECHANISM IS DISTRIBUTED IN SUCH NTCB SUBJECTS (I.E., WHEN OUTSIDE THE REFERENCE MONITOR), THE ASSURANCE REQUIREMENTS (SEE THE ASSURANCE SECTION) FOR THE DESIGN AND IMPLEMENTATION OF THE DAC SHALL BE THOSE OF CLASS C2 FOR ALL NETWORKS OF CLASS C2 OR ABOVE.

When integrity is included as part of the network discretionary security policy, the above interpretations shall be specifically applied to the controls over modification, viz, the write mode of access, within each component based on identified users or groups of users.

+ Rationale

In this class, THE SUPPORTING ELEMENTS OF THE OVERALL DAC MECHANISM ARE REQUIRED TO ISOLATE INFORMATION (OBJECTS) THAT SUPPORTS DAC SO THAT IT IS SUBJECT TO AUDITING REQUIREMENTS (SEE THE SYSTEM ARCHITECTURE SECTION). THE USE OF NETWORK IDENTIFIERS TO IDENTIFY GROUPS OF INDIVIDUAL USERS COULD BE IMPLEMENTED, FOR EXAMPLE, AS AN X.25 COMMUNITY OF INTEREST IN THE NETWORK PROTOCOL LAYER (LAYER 3). IN ALL OTHER RESPECTS, the supporting elements of the overall DAC mechanism are treated exactly as untrusted subjects are treated with respect to DAC in an ADP system, with the same

result as noted in the interpretation.

A typical situation for DAC is that a surrogate process for a remote user will be created in some host for access to objects under the control of the NTCB partition within that host. The interpretation requires that a user identifier be assigned and maintained for each such process by the NTCB, so that access by a surrogate process is subject to essentially the same discretionary controls as access by a process acting on behalf of a local user would be. However, within this interpretation a range of possible interpretations of the assigned user identification is permitted.

The most obvious situation would exist if a global database of network users were to be made permanently available on demand to every host, (i.e., a name server existed) so that all user identifications were globally meaningful.

It is also acceptable, however, for some NTCB partitions to maintain a database of locally-registered users for its own use. In such a case, one could choose to inhibit the creation of surrogate processes for locally unregistered users, or (if permitted by the local policy) alternatively, to permit the creation of surrogate processes with preselected user and group identifiers which, in effect, identify the process as executing on behalf of a member of a group of users on a particular remote host. The intent of the words concerning audit in the interpretation is to provide a minimally acceptable degree of auditability for cases such as the last described. What is required is that there be a capability, using the audit facilities provided by the network NTCB partitions involved, to determine who was logged in at the actual host of the group of remote users at the time the surrogate processing occurred.

Associating the proper user id with a surrogate process is the job of identification and authentication. This means that DAC is applied locally, with respect to the user id of the surrogate process. The transmission of the data back across the network to the user's host, and the creation of a copy of the data there, is not the business of DAC.

Components that support only internal subjects impact the implementation of the DAC by providing services by which information (e.g., a user-id) is made available to a component that makes a DAC decision. An example of the latter would be the case that a user at Host A attempts to access a file at Host B. The DAC decision might be (and usually would be) made at Host B on the basis of a user-id transmitted from Host A to Host B.

Unique user identification may be achieved by a variety of mechanisms, including (a) a requirement for unique identification and authentication on the host where access takes

place; (b) recognition of fully qualified network addresses authenticated by another host and forwarded to the host where access takes place; or (c) administrative support of a network-wide unique personnel identifier that could be authenticated and forwarded by another host as in (b) above, or could be authenticated and forwarded by a dedicated network identification and authentication server. The protocols which implement (b) or (c) are subject to the System Architecture requirements.

Network support for DAC might be handled in other ways than that described as "typical" above. In particular, some form of centralized access control is often proposed. An access control center may make all decisions for DAC, or it may share the burden with the hosts by controlling host-to-host connections, and leaving the hosts to decide on access to their objects by users at a limited set of remote hosts. In this case the access control center provides the linkage between the connection oriented abstraction (as discussed in the Introduction) and the overall network security policy for DAC. In all cases the enforcement of the decision must be provided by the host where the object resides.

2.2.1.2 Object Reuse

+ Statement from DoD 5200.28-STD

ALL AUTHORIZATIONS TO THE INFORMATION CONTAINED WITHIN A STORAGE OBJECT SHALL BE REVOKED PRIOR TO INITIAL ASSIGNMENT, ALLOCATION OR REALLOCATION TO A SUBJECT FROM THE TCB'S POOL OF UNUSED STORAGE OBJECTS. NO INFORMATION, INCLUDING ENCRYPTED REPRESENTATIONS OF INFORMATION, PRODUCED BY A PRIOR SUBJECT'S ACTIONS IS TO BE AVAILABLE TO ANY SUBJECT THAT OBTAINS ACCESS TO AN OBJECT THAT HAS BEEN RELEASED BACK TO THE SYSTEM.

+ Interpretation

THE NTCB SHALL ENSURE THAT ANY STORAGE OBJECTS THAT IT CONTROLS (E.G., MESSAGE BUFFERS UNDER THE CONTROL OF A NTCB PARTITION IN A COMPONENT) CONTAIN NO INFORMATION FOR WHICH A SUBJECT IN THAT COMPONENT IS NOT AUTHORIZED BEFORE GRANTING ACCESS. THIS REQUIREMENT MUST BE ENFORCED BY EACH OF THE NTCB PARTITIONS.

+ Rationale

IN A NETWORK SYSTEM, STORAGE OBJECTS OF INTEREST ARE THINGS THAT THE NTCB DIRECTLY CONTROLS, SUCH AS MESSAGE BUFFERS IN COMPONENTS. EACH COMPONENT OF THE NETWORK SYSTEM MUST ENFORCE THE OBJECT REUSE REQUIREMENT WITH RESPECT TO THE STORAGE OBJECTS OF INTEREST AS DETERMINED BY THE NETWORK SECURITY POLICY. FOR EXAMPLE, THE DAC REQUIREMENT IN THIS DIVISION LEADS TO THE REQUIREMENT HERE THAT MESSAGE BUFFERS

BE UNDER THE CONTROL OF THE NTCB PARTITION. A BUFFER ASSIGNED TO AN INTERNAL SUBJECT MAY BE REUSED AT THE DISCRETION OF THAT SUBJECT WHICH IS RESPONSIBLE FOR PRESERVING THE INTEGRITY OF MESSAGE STREAMS. SUCH CONTROLLED OBJECTS MAY BE IMPLEMENTED IN PHYSICAL RESOURCES, SUCH AS BUFFERS, DISK SECTORS, TAPE SPACE, AND MAIN MEMORY, IN COMPONENTS SUCH AS NETWORK SWITCHES.

2.2.2 Accountability

2.2.2.1 Identification and Authentication

+ Statement from DoD 5200.28-STD

The TCB shall require users to identify themselves to it before beginning to perform any other actions that the TCB is expected to mediate. Furthermore, the TCB shall use a protected mechanism (e.g., passwords) to authenticate the user's identity. The TCB shall protect authentication data so that it cannot be accessed by any unauthorized user. THE TCB SHALL BE ABLE TO ENFORCE INDIVIDUAL ACCOUNTABILITY BY PROVIDING THE CAPABILITY TO UNIQUELY IDENTIFY EACH INDIVIDUAL ADP SYSTEM USER. THE TCB SHALL ALSO PROVIDE THE CAPABILITY OF ASSOCIATING THIS IDENTITY WITH ALL AUDITABLE ACTIONS TAKEN BY THAT INDIVIDUAL.

+ Interpretation

The requirement for identification and authentication of users is the same for a network system as for an ADP system. The identification and authentication may be done by the component to which the user is directly connected or some other component, such as an identification and authentication server. Available techniques, such as those described in the Password Guideline, are generally also applicable in the network context. However, in cases where the NTCB is expected to mediate actions of a host (or other network component) that is acting on behalf of a user or group of users, the NTCB may employ identification and authentication of the host (or other component) in lieu of identification and authentication of an individual user, SO LONG AS THE COMPONENT IDENTIFIER IMPLIES A LIST OF SPECIFIC USERS UNIQUELY ASSOCIATED WITH THE IDENTIFIER AT THE TIME OF ITS USE FOR AUTHENTICATION. THIS REQUIREMENT DOES NOT APPLY TO INTERNAL SUBJECTS.

Authentication information, including the identity of a user (once authenticated) may be passed from one component to another without reauthentication, so long as the NTCB protects (e.g., by encryption) the information from unauthorized disclosure and modification. This protection shall

provide at least a similar level of assurance (or strength of mechanism) as pertains to the protection of the authentication mechanism and authentication data.

+ Rationale

The need for accountability is not changed in the context of a network system. The fact that the NTCB is partitioned over a set of components neither reduces the need nor imposes new requirements. That is, individual accountability is still the objective. ALSO, in the context of a network system at the (C2) LEVEL OR HIGHER "INDIVIDUAL ACCOUNTABILITY" CAN BE SATISFIED BY IDENTIFICATION OF A HOST (OR OTHER COMPONENT) SO LONG AS THE REQUIREMENT FOR TRACEABILITY TO INDIVIDUAL USERS OR A SET OF SPECIFIC INDIVIDUAL USERS WITH ACTIVE SUBJECTS IS SATISFIED. THERE IS ALLOWED TO BE AN UNCERTAINTY IN TRACEABILITY BECAUSE OF ELAPSED TIME BETWEEN CHANGES IN THE GROUP MEMBERSHIP AND THE ENFORCEMENT IN THE ACCESS CONTROL MECHANISMS. In addition, there is no need in a distributed processing system like a network to reauthenticate a user at each point in the network where a projection of a user (via the subject operating on behalf of the user) into another remote subject takes place.

= Department of Defense Password Management Guide-
line, CSC-STD-002-85

The passing of identifiers and/or authentication information from one component to another is usually done in support to the implementation of the discretionary access control (DAC). This support relates directly to the DAC regarding access by a user to a storage object in a different NTCB partition than the one where the user was authenticated. Employing a forwarded identification implies additional reliance on the source and components along the path.

2.2.2.2 Audit

+ Statement from DoD 5200.28-STD

THE TCB SHALL BE ABLE TO CREATE, MAINTAIN, AND PROTECT FROM MODIFICATION OR UNAUTHORIZED ACCESS OR DESTRUCTION AN AUDIT TRAIL OF ACCESSES TO THE OBJECTS IT PROTECTS. THE AUDIT DATA SHALL BE PROTECTED BY THE TCB SO THAT READ ACCESS TO IT IS LIMITED TO THOSE WHO ARE AUTHORIZED FOR AUDIT DATA. THE TCB SHALL BE ABLE TO RECORD THE FOLLOWING TYPES OF EVENTS: USE OF IDENTIFICATION AND AUTHENTICATION MECHANISMS, INTRODUCTION OF OBJECTS INTO A USER'S ADDRESS SPACE (E.G., FILE OPEN, PROGRAM INITIATION), DELETION OF OBJECTS, ACTIONS TAKEN BY COMPUTER OPERATORS AND SYSTEM ADMINISTRATORS AND/OR

SYSTEM SECURITY OFFICERS, AND OTHER SECURITY RELEVANT EVENTS. FOR EACH RECORDED EVENT, THE AUDIT RECORD SHALL IDENTIFY: DATE AND TIME OF THE EVENT, USER, TYPE OF EVENT, AND SUCCESS OR FAILURE OF THE EVENT. FOR IDENTIFICATION/AUTHENTICATION EVENTS THE ORIGIN OF REQUEST (E.G., TERMINAL ID) SHALL BE INCLUDED IN THE AUDIT RECORD. FOR EVENTS THAT INTRODUCE AN OBJECT INTO A USER'S ADDRESS SPACE AND FOR OBJECT DELETION EVENTS THE AUDIT RECORD SHALL INCLUDE THE NAME OF THE OBJECT. THE ADP SYSTEM ADMINISTRATOR SHALL BE ABLE TO SELECTIVELY AUDIT THE ACTIONS OF ANY ONE OR MORE USERS BASED ON INDIVIDUAL IDENTITY.

+ Interpretation

THIS CRITERION APPLIES AS STATED. THE SPONSOR MUST SELECT WHICH EVENTS ARE AUDITABLE. IF ANY SUCH EVENTS ARE NOT DISTINGUISHABLE BY THE NTCB ALONE (FOR EXAMPLE THOSE IDENTIFIED IN PART II), THE AUDIT MECHANISM SHALL PROVIDE AN INTERFACE, WHICH AN AUTHORIZED SUBJECT CAN INVOKE WITH PARAMETERS SUFFICIENT TO PRODUCE AN AUDIT RECORD. THESE AUDIT RECORDS SHALL BE DISTINGUISHABLE FROM THOSE PROVIDED BY THE NTCB. IN THE CONTEXT OF A NETWORK SYSTEM, "OTHER SECURITY RELEVANT EVENTS" (DEPENDING ON NETWORK SYSTEM ARCHITECTURE AND NETWORK SECURITY POLICY) MIGHT BE AS FOLLOWS:

1. IDENTIFICATION OF EACH ACCESS EVENT (E.G., ESTABLISHING A CONNECTION OR A CONNECTIONLESS ASSOCIATION BETWEEN PROCESSES IN TWO HOSTS OF THE NETWORK) AND ITS PRINCIPAL PARAMETERS (E.G., HOST IDENTIFIERS OF THE TWO HOSTS INVOLVED IN THE ACCESS EVENT AND USER IDENTIFIER OR HOST IDENTIFIER OF THE USER OR HOST THAT IS REQUESTING THE ACCESS EVENT)
2. IDENTIFICATION OF THE STARTING AND ENDING TIMES OF EACH ACCESS EVENT USING LOCAL TIME OR GLOBAL SYNCHRONIZED TIME
3. IDENTIFICATION OF SECURITY-RELEVANT EXCEPTIONAL CONDITIONS (E.G., POTENTIAL VIOLATION OF DATA INTEGRITY, SUCH AS MISROUTED DATAGRAMS) DETECTED DURING THE TRANSACTIONS BETWEEN TWO HOSTS
4. UTILIZATION OF CRYPTOGRAPHIC VARIABLES
5. CHANGING THE CONFIGURATION OF THE NETWORK (E.G., A COMPONENT LEAVING THE NETWORK AND REJOINING)

IN ADDITION, IDENTIFICATION INFORMATION SHOULD BE INCLUDED IN APPROPRIATE AUDIT TRAIL RECORDS, AS NECESSARY, TO ALLOW ASSOCIATION OF ALL RELATED (E.G., INVOLVING THE SAME NETWORK EVENT) AUDIT TRAIL RECORDS (E.G., AT DIFFERENT HOSTS) WITH EACH OTHER. FURTHERMORE, A COMPONENT OF THE NETWORK SYSTEM MAY PROVIDE THE REQUIRED AUDIT CAPABILITY

(E.G., STORAGE, RETRIEVAL, REDUCTION, ANALYSIS) FOR OTHER COMPONENTS THAT DO NOT INTERNALLY STORE AUDIT DATA BUT TRANSMIT THE AUDIT DATA TO SOME DESIGNATED COLLECTION COMPONENT. PROVISIONS SHALL BE MADE TO CONTROL THE LOSS OF AUDIT DATA DUE TO UNAVAILABILITY OF RESOURCES.

IN THE CONTEXT OF A NETWORK SYSTEM, THE "USER'S ADDRESS SPACE" IS EXTENDED, FOR OBJECT INTRODUCTION AND DELETION EVENTS, TO INCLUDE ADDRESS SPACES BEING EMPLOYED ON BEHALF OF A REMOTE USER (OR HOST). HOWEVER, THE FOCUS REMAINS ON USERS IN CONTRAST TO INTERNAL SUBJECTS AS DISCUSSED IN THE DAC CRITERION. IN ADDITION, AUDIT INFORMATION MUST BE STORED IN MACHINE-READABLE FORM.

+ Rationale

FOR REMOTE USERS, THE NETWORK IDENTIFIERS (E.G., INTERNET ADDRESS) CAN BE USED AS IDENTIFIERS OF GROUPS OF INDIVIDUAL USERS (E.G., "ALL USERS AT HOST A") TO ELIMINATE THE MAINTENANCE THAT WOULD BE REQUIRED IF INDIVIDUAL IDENTIFICATION OF REMOTE USERS WAS EMPLOYED. IN THIS CLASS (C2), HOWEVER, IT MUST BE POSSIBLE TO IDENTIFY (IMMEDIATELY OR AT SOME LATER TIME) THE INDIVIDUALS REPRESENTED BY A GROUP IDENTIFIER. IN ALL OTHER RESPECTS, THE INTERPRETATION IS A STRAIGHTFORWARD EXTENSION OF THE CRITERION INTO THE CONTEXT OF A NETWORK SYSTEM.

2.2.3 Assurance

— — — — —

2.2.3.1 Operational Assurance

2.2.3.1.1 System Architecture

+ Statement from DoD 5200.28-STD

The TCB shall maintain a domain for its own execution that protects it from external interference or tampering (e.g., by modification of its code or data structures). Resources controlled by the TCB may be a defined subset of the subjects and objects in the ADP system. THE TCB SHALL ISOLATE THE RESOURCES TO BE PROTECTED SO THAT THEY ARE SUBJECT TO THE ACCESS CONTROL AND AUDITING REQUIREMENTS.

+ Interpretation

The system architecture criterion must be met individually by all NTCB partitions. Implementation of the requirement that the NTCB maintain a domain for its own execution is achieved by having each NTCB partition maintain a domain for its own execution.

The subset of network resources over which the NTCB has control are the union of the sets of resources over which the NTCB partitions have control. Code and data structures belonging to the NTCB, transferred among NTCB subjects (i.e., subjects outside the reference monitor but inside the NTCB) belonging to different NTCB partitions, must be protected against external interference or tampering. For example, a cryptographic checksum or physical means may be employed to protect user authentication data exchanged between NTCB partitions.

EACH NTCB PARTITION PROVIDES ISOLATION OF RESOURCES (WITHIN ITS COMPONENT) TO BE PROTECTED IN ACCORD WITH THE NETWORK SYSTEM ARCHITECTURE AND SECURITY POLICY.

+ Rationale

The requirement for the protection of communications between NTCB partitions is specifically directed to subjects that are part of the NTCB partitions. Any requirements for such protection for the subjects that are outside the NTCB partitions are addressed in response to the integrity requirements of the security policy.

ISOLATION OF THE RESOURCES TO BE PROTECTED PROVIDES ADDITIONAL PROTECTION, COMPARED TO CLASS (C1), THAT MECHANISMS THAT DEPEND ON THE RESOURCE (E.G., DAC AND USER IDENTIFICATION) WILL OPERATE CORRECTLY.

2.2.3.1.2 System Integrity

+ Statement from DoD 5200.28-STD

Hardware and/or software features shall be provided that can be used to periodically validate the correct operation of the on-site hardware and firmware elements of the TCB.

+ Interpretation

Implementation of the requirement is partly achieved by having hardware and/or software features that can be used to periodically validate the correct operation of the hardware and firmware elements of each component's NTCB partition. Features shall also be provided to validate the identity and correct operation of a component prior to its incorporation in the network system and throughout system operation. For example, a protocol could be designed that enables the components of the partitioned NTCB to exchange messages periodically and validate each other's correct response. The protocol shall be able to determine the remote entity's ability to respond. NTCB partitions shall provide the capability to report to network administrative personnel the failures

detected in other NTCB partitions.

Intercomponent protocols implemented within a NTCB shall be designed in such a way as to provide correct operation in the case of failures of network communications or individual components. The allocation of discretionary access control policy in a network may require communication between trusted subjects that are part of the NTCB partitions in different components. This communication is normally implemented with a protocol between the subjects as peer entities. Incorrect access within a component shall not result from failure of an NTCB partition to communicate with other components.

+ Rationale

The first paragraph of the interpretation is a straightforward extension of the requirement into the context of a network system and partitioned NTCB as defined for these network criteria.

NTCB protocols should be robust enough so that they permit the system to operate correctly in the case of localized failure. The purpose of this protection is to preserve the integrity of the NTCB itself. It is not unusual for one or more components in a network to be inoperative at any time, so it is important to minimize the effects of such failures on the rest of the network. Additional integrity and denial of service issues are addressed in Part II.

2.2.3.2 Life-Cycle Assurance

2.2.3.2.1 Security Testing

+ Statement from DoD 5200.28-STD

The security mechanisms of the ADP system shall be tested and found to work as claimed in the system documentation. Testing shall be done to assure that there are no obvious ways for an unauthorized user to bypass or otherwise defeat the security protection mechanisms of the TCB. TESTING SHALL ALSO INCLUDE A SEARCH FOR OBVIOUS FLAWS THAT WOULD ALLOW VIOLATION OF RESOURCE ISOLATION, OR THAT WOULD PERMIT UNAUTHORIZED ACCESS TO THE AUDIT OR AUTHENTICATION DATA. (See the Security Testing Guidelines.)

+ Interpretation

Testing of a component will require a testbed that exercises the interfaces and protocols of the COMPONENT INCLUDING TESTS UNDER EXCEPTIONAL CONDITIONS. The testing of a security mechanism of the network system for meeting this criterion shall be an integrated testing procedure

involving all components containing an NTCB partition that implement the given mechanism. This integrated testing is additional to any individual component tests involved in the evaluation of the network system. The sponsor should identify the allowable set of configurations including the sizes of the networks. Analysis or testing procedures and tools shall be available to test the limits of these configurations. A change in configuration within the allowable set of configurations does not require retesting.

+ Rationale

Testing is the primary method available in this evaluation division to gain any assurance that the security mechanisms perform their intended function.

2.2.4 Documentation

2.2.4.1 Security Features User's Guide

+ Statement from DoD 5200.28-STD

A single summary, chapter, or manual in user documentation shall describe the protection mechanisms provided by the TCB, interpretations on their use, and how they interact with one another.

+ Interpretation

This user documentation describes user visible protection mechanisms at the global (network system) level and at the user interface of each component, and the interaction among these.

+ Rationale

The interpretation is an extension of the requirement into the context of a network system as defined for these network criteria. Documentation of protection mechanisms provided by individual components is required by the criteria for trusted computer systems that are applied as appropriate for the individual components.

2.2.4.2 Trusted Facility Manual

+ Statement from DoD 5200.28-STD

A manual addressed to the ADP system administrator shall present cautions about functions and privileges that should be controlled when running a secure facility. THE PROCEDURES

FOR EXAMINING AND MAINTAINING THE AUDIT FILES AS WELL AS THE DETAILED AUDIT RECORD STRUCTURE FOR EACH TYPE OF AUDIT EVENT SHALL BE GIVEN.

+ Interpretation

This manual shall contain specifications and procedures to assist the system administrator(s) maintain cognizance of the network configuration. These specifications and procedures shall address the following:

1. The hardware configuration of the network itself;
2. The implications of attaching new components to the network;
3. The case where certain components may periodically leave the network (e.g., by crashing, or by being disconnected) and then rejoin;
4. Network configuration aspects that can impact the security of the network system; (For example, the manual should describe for the network system administrator the interconnections among components that are consistent with the overall network system architecture.)
5. Loading or modifying NTCB software or firmware (e.g., down-line loading).

The physical and administrative environmental controls shall be specified. Any assumptions about security of a given network should be clearly stated (e.g., the fact that all communications links must be physically protected to a certain level).

+ Rationale

There may be multiple system administrators with diverse responsibilities. The technical security measures described by these criteria must be used in conjunction with other forms of security in order to achieve security of the network. Additional forms include administrative security, physical security, emanations security, etc.

Extension of this criterion to cover configuration aspects of the network is needed because, for example, proper interconnection of components is typically essential to achieve a correct realization of the network architecture.

Cryptography is one common mechanism employed to protect communication circuits. Encryption transforms the representation of information so that it is unintelligible

to unauthorized subjects. Reflecting this transformation, the sensitivity of the ciphertext is generally lower than the cleartext. If encryption methodologies are employed, they shall be approved by the National Security Agency (NSA).

The encryption algorithm and its implementation are outside the scope of these interpretations. This algorithm and implementation may be implemented in a separate device or may be a function of a subject in a component not dedicated to encryption. Without prejudice, either implementation packaging is referred to as an encryption mechanism herein.

2.2.4.3 Test Documentation

+ Statement from DoD 5200.28-STD

The system developer shall provide to the evaluators a document that describes the test plan, test procedures that show how the security mechanisms were tested, and results of the security mechanisms' functional testing.

+ Interpretation

The "system developer" is interpreted as "the network system sponsor". The description of the test plan should establish the context in which the testing was or should be conducted. The description should identify any additional test components that are not part of the system being evaluated. This includes a description of the test-relevant functions of such test components and a description of the interfacing of those test components to the system being evaluated. The description of the test plan should also demonstrate that the tests adequately cover the network security policy. The tests should include the features described in the System Architecture and the System Integrity sections. The tests should also include network configuration and sizing.

+ Rationale

The entity being evaluated may be a networking subsystem (see Appendix A) to which other components must be added to make a complete network system. In that case, this interpretation is extended to include contextual definition because, at evaluation time, it is not possible to validate the test plans without the description of the context for testing the networking subsystem.

2.2.4.4 Design Documentation

+ Statement from DoD 5200.28-STD

Documentation shall be available that provides a description of the manufacturer's philosophy of protection and an explanation of how this philosophy is translated into the TCB. If the TCB is composed of distinct modules, the interfaces between these modules shall be described.

+ Interpretation

Explanation of how the sponsor's philosophy of protection is translated into the NTCB shall include a description of how the NTCB is partitioned. The security policy also shall be stated. The description of the interfaces between the NTCB modules shall include the interface(s) between NTCB partitions and modules within the partitions if the modules exist. The sponsor shall describe the security architecture and design, including the allocation of security requirements among components. Appendix A addresses component evaluation issues.

+ Rationale

The interpretation is a straightforward extension of the requirement into the context of a network system as defined for this network interpretation. Other documentation, such as description of components and description of operating environment(s) in which the networking subsystem or network system is designed to function, is required elsewhere, e.g., in the Trusted Facility Manual.

In order to be evaluated, a network must possess a coherent Network Security Architecture and Design. (Interconnection of components that do not adhere to such a single coherent Network Security Architecture is addressed in the Interconnection of Accredited AIS, Appendix C.) The Network Security Architecture must address the security-relevant policies, objectives, and protocols. The Network Security Design specifies the interfaces and services that must be incorporated into the network so that it can be evaluated as a trusted entity. There may be multiple designs that conform to the same architecture but are more or less incompatible and non-interoperable (except through the Interconnection Rules). Security related mechanisms requiring cooperation among components are specified in the design in terms of their visible interfaces; mechanisms having no visible interfaces are not specified in this document but are left as implementation decisions.

The Network Security Architecture and Design must be available from the network sponsor before evaluation of the network, or any component, can be undertaken. The Network Security Architecture and Design must be sufficiently complete, unambiguous, and free from obvious flaws to permit

the construction or assembly of a trusted network based on the structure it specifies.

When a component is being designed or presented for evaluation, or when a network assembled from components is assembled or presented for evaluation, there must be a priori evidence that the Network security Architecture and Design are satisfied. That is, the components can be assembled into a network that conforms in every way with the Network Security Architecture and Design to produce a physical realization that is trusted to the extent that its evaluation indicates.

In order for a trusted network to be constructed from components that can be built independently, the Network Security Architecture and Design must completely and unambiguously define the security functionality of components as well as the interfaces between or among components. The Network Security Architecture and Design must be evaluated to determine that a network constructed to its specifications will in fact be trusted, that is, it will be evaluable under these interpretations.

3.0 DIVISION B: MANDATORY PROTECTION

The notion of an NTCB that preserves the integrity of sensitivity labels and uses them to enforce a set of mandatory access control rules is a major requirement in this division. Network systems in this division must carry the sensitivity labels with major data structures in the system. The network system sponsor also provides the security policy model on which the NTCB is based and furnishes a specification of the NTCB. Evidence must be provided to demonstrate that the reference monitor concept has been implemented.

3.1 CLASS (B1): LABELED SECURITY PROTECTION

CLASS (B1) NETWORK SYSTEMS REQUIRE ALL THE FEATURES REQUIRED FOR CLASS (C2). IN ADDITION, AN INFORMAL STATEMENT OF THE SECURITY POLICY MODEL, DATA LABELING, AND MANDATORY ACCESS CONTROL OVER SUBJECTS AND STORAGE OBJECTS MUST BE PRESENT. THE CAPABILITY MUST EXIST FOR ACCURATELY LABELING EXPORTED INFORMATION. ANY FLAWS IDENTIFIED BY TESTING MUST BE REMOVED. THE FOLLOWING ARE MINIMAL REQUIREMENTS FOR NETWORK SYSTEMS ASSIGNED A CLASS (B1) RATING:

3.1.1 Security Policy

+ Statement from DoD 5200.28-STD

Implied from the Introduction to the TCSEC.

+ Interpretation

The network sponsor shall describe the overall network security policy enforced by the NTCB. At a minimum, this policy shall include the discretionary AND MANDATORY requirements applicable to this class. The policy may require data secrecy, or data integrity, or both. THE POLICY IS AN ACCESS CONTROL POLICY HAVING TWO PRIMARY COMPONENTS: MANDATORY AND DISCRETIONARY. The policy shall include a discretionary policy for protecting the information being processed based on the authorizations of individuals, users, or groups of users. This access control policy statement shall describe the requirements on the network to prevent or detect "reading or destroying" sensitive information by unauthorized users or errors. THE MANDATORY POLICY MUST DEFINE THE SET OF DISTINCT SENSITIVITY LEVELS THAT IT SUPPORTS. FOR THE CLASS B1 OR ABOVE THE MANDATORY POLICY SHALL BE BASED ON THE LABELS ASSOCIATED WITH THE INFORMATION THAT REFLECTS ITS SENSITIVITY WITH RESPECT TO SECRECY AND/OR INTEGRITY, WHERE APPLICABLE, AND LABELS ASSOCIATED WITH USERS TO REFLECT THEIR AUTHORIZATION TO ACCESS SUCH INFORMATION. UNAUTHORIZED USERS INCLUDE BOTH THOSE that are not authorized to use the network at all (e.g., a user attempting to use a passive or active wire tap) or a legitimate user of the network who is not authorized to access a specific piece of information being protected.

Note that "users" does not include "operators," "system programmers," "technical control officers," "system security officers," and other system support personnel. They are distinct from users and are subject to the Trusted Facility Manual and the System Architecture requirements. Such individuals may change the system parameters of the network system, for example, by defining membership of a group. These individuals may also have the separate role of users.

SECRECY POLICY: The network sponsor shall define the form of the discretionary AND MANDATORY secrecy policy that is enforced in the network to prevent unauthorized users from reading the sensitive information entrusted to the network.

DATA INTEGRITY POLICY: The network sponsor shall define the discretionary AND MANDATORY integrity policy to prevent unauthorized users from modifying,

viz., writing, sensitive information. The definition of data integrity presented by the network sponsor refers to the requirement that the information has not been subjected to unauthorized modification in the network. THE MANDATORY INTEGRITY POLICY ENFORCED BY THE NTCB CANNOT, IN GENERAL, PREVENT MODIFICATION WHILE INFORMATION IS BEING TRANSMITTED BETWEEN COMPONENTS. HOWEVER, AN INTEGRITY SENSITIVITY LABEL MAY REFLECT THE CONFIDENCE THAT THE INFORMATION HAS NOT BEEN SUBJECTED TO TRANSMISSION ERRORS BECAUSE OF THE PROTECTION AFFORDED DURING TRANSMISSION. THIS REQUIREMENT IS DISTINCT FROM THE REQUIREMENT FOR LABEL INTEGRITY.

+ Rationale

The word "sponsor" is used in place of alternatives (such as "vendor," "architect," "manufacturer," and "developer") because the alternatives indicate people who may not be available, involved, or relevant at the time that a network system is proposed for evaluation.

A trusted network is able to control both the reading and writing of shared sensitive information. Control of writing is used to protect against destruction of information. A network normally is expected to have policy requirements to protect both the secrecy and integrity of the information entrusted to it. In a network the integrity is frequently as important or more important than the secrecy requirements. Therefore the secrecy and/or integrity policy to be enforced by the network must be stated for each network regardless of its evaluation class. The assurance that the policy is faithfully enforced is reflected in the evaluation class of the network.

This control over modification is typically used to protect information so that it may be relied upon and to control the potential harm that would result if the information were corrupted. The overall network policy requirements for integrity includes the protection for data both while being processed in a component and while being transmitted in the network. The access control policy enforced by the NTCB relates to the access of subjects to objects within each component. Communications integrity addressed within Part II relates to information while being transmitted.

THE MANDATORY INTEGRITY POLICY (AT CLASS B1 AND ABOVE) IN SOME ARCHITECTURES MAY BE USEFUL IN SUPPORTING THE LINKAGE BETWEEN THE CONNECTION ORIENTED ABSTRACTION INTRODUCED IN THE INTRODUCTION AND THE INDIVIDUAL COMPONENTS OF THE NETWORK. FOR EXAMPLE, IN A KEY DISTRIBUTION CENTER FOR END-TO-END ENCRYPTION, A DISTINCT INTEGRITY CATEGORY MAY BE ASSIGNED TO ISOLATE THE KEY GENERATION CODE AND DATA FROM

POSSIBLE MODIFICATION BY OTHER SUPPORTING PROCESSES IN THE SAME COMPONENT, SUCH AS OPERATOR INTERFACES AND AUDIT.

THE MANDATORY INTEGRITY POLICY FOR SOME ARCHITECTURE MAY DEFINE AN INTEGRITY SENSITIVITY LABEL THAT REFLECTS THE SPECIFIC REQUIREMENTS FOR ENSURING THAT INFORMATION HAS NOT BEEN SUBJECT TO RANDOM ERRORS IN EXCESS OF A STATED LIMIT NOR TO UNAUTHORIZED MESSAGE STREAM MODIFICATION (MSM) -. THE SPECIFIC METRIC ASSOCIATED WITH AN INTEGRITY SENSITIVITY LABEL WILL GENERALLY REFLECT THE INTENDED APPLICATIONS OF THE NETWORK.

3.1.1.1 Discretionary Access Control

+ Statement from DoD 5200.28-STD

The TCB shall define and control access between named users and named objects (e.g., files and programs) in the ADP system. The enforcement mechanism (e.g., self/group/public controls, access control lists) shall allow users to specify and control sharing of those objects by named individuals or defined groups of individuals, or both, and shall provide controls to limit propagation of access rights. The discretionary access control mechanism shall, either by explicit user action or by default, provide that objects are protected from unauthorized access. These access controls shall be capable of including or excluding access to the granularity of a single user. Access permission to an object by users not already possessing access permission shall only be assigned by authorized users.

- See Voydock, Victor L. and Stephen T. Kent, "Security Mechanisms in High-Level Network Protocols," Computing Surveys, Vol. 15, No. 2, June 1983, pp 135-171.

+ Interpretation

The discretionary access control (DAC) mechanism(s) may be distributed over the partitioned NTCB in various ways. Some part, all, or none of the DAC may be implemented in a given component of the network system. In particular, components that support only internal subjects (i.e., that have no subjects acting as direct surrogates for users), such as a public network packet switch, might not implement the DAC mechanism(s) directly (e.g., they are unlikely to contain access control lists).

Identification of users by groups may be achieved in various ways in the networking environment. For example, the network identifiers (e.g., internet addresses) for various components (e.g., hosts, gateways) can be used as iden-

tifiers of groups of individual users (e.g., "all users at Host A," "all users of network Q") so long as the individuals involved in the group are implied by the group identifier. For example, Host A might employ a particular group-id, for which it maintains a list of explicit users in that group, in its network exchange with Host B, which accepts the group-id under the conditions of this interpretation.

For networks, individual hosts will impose need-to-know controls over their users on the basis of named individuals - much like (in fact, probably the same) controls used when there is no network connection.

When group identifiers are acceptable for access control, the identifier of some other host may be employed, to eliminate the maintenance that would be required if individual identification of remote users was employed. In class C2 and higher, however, it must be possible from that audit record to identify (immediately or at some later time) exactly the individuals represented by a group identifier at the time of the use of that identifier. There is allowed to be an uncertainty because of elapsed time between changes in the group membership and the enforcement in the access control mechanisms.

The DAC mechanism of a NTCB partition may be implemented at the interface of the reference monitor or may be distributed in subjects that are part of the NTCB in the same or different component. The reference monitor manages all the physical resources of the system and from them creates the abstraction of subjects and objects that it controls. Some of these subjects and objects may be used to implement a part of the NTCB. When the DAC mechanism is distributed in such NTCB subjects (i.e., when outside the reference monitor), the assurance requirements (see the Assurance section) for the design and implementation of the DAC shall be those of class C2 for all networks of class C2 or above.

When integrity is included as part of the network discretionary security policy, the above interpretations shall be specifically applied to the controls over modification, viz, the write mode of access, within each component based on identified users or groups of users.

+ Rationale

In this class, the supporting elements of the overall DAC mechanism are required to isolate information (objects) that supports DAC so that it is subject to auditing requirements (see the System Architecture section). The use of network identifiers to identify groups of individual users could be implemented, for example, as an X.25 community of interest in the network protocol layer (layer 3). In all

other respects, the supporting elements of the overall DAC mechanism are treated exactly as untrusted subjects are treated with respect to DAC in an ADP system, with the same result as noted in the interpretation.

A typical situation for DAC is that a surrogate process for a remote user will be created in some host for access to objects under the control of the NTCB partition within that host. The interpretation requires that a user identifier be assigned and maintained for each such process by the NTCB, so that access by a surrogate process is subject to essentially the same discretionary controls as access by a process acting on behalf of a local user would be. However, within this interpretation a range of possible interpretations of the assigned user identification is permitted.

The most obvious situation would exist if a global database of network users were to be made permanently available on demand to every host, (i.e., a name server existed) so that all user identifications were globally meaningful.

It is also acceptable, however, for some NTCB partitions to maintain a database of locally-registered users for its own use. In such a case, one could choose to inhibit the creation of surrogate processes for locally unregistered users, or (if permitted by the local policy) alternatively, to permit the creation of surrogate processes with preselected user and group identifiers which, in effect, identify the process as executing on behalf of a member of a group of users on a particular remote host. The intent of the words concerning audit in the interpretation is to provide a minimally acceptable degree of auditability for cases such as the last described. What is required is that there be a capability, using the audit facilities provided by the network NTCB partitions involved, to determine who was logged in at the actual host of the group of remote users at the time the surrogate processing occurred.

Associating the proper user id with a surrogate process is the job of identification and authentication. This means that DAC is applied locally, with respect to the user id of the surrogate process. The transmission of the data back across the network to the user's host, and the creation of a copy of the data there, is not the business of DAC.

Components that support only internal subjects impact the implementation of the DAC by providing services by which information (e.g., a user-id) is made available to a component that makes a DAC decision. An example of the latter would be the case that a user at Host A attempts to access a file at Host B. The DAC decision might be (and usually would be) made at Host B on the basis of a user-id transmitted from Host A to Host B.

Unique user identification may be achieved by a variety of mechanisms, including (a) a requirement for unique identification and authentication on the host where access takes place; (b) recognition of fully qualified network addresses authenticated by another host and forwarded to the host where access takes place; or (c) administrative support of a network-wide unique personnel identifier that could be authenticated and forwarded by another host as in (b) above, or could be authenticated and forwarded by a dedicated network identification and authentication server. The protocols which implement (b) or (c) are subject to the System Architecture requirements.

Network support for DAC might be handled in other ways than that described as "typical" above. In particular, some form of centralized access control is often proposed. An access control center may make all decisions for DAC, or it may share the burden with the hosts by controlling host-to-host connections, and leaving the hosts to decide on access to their objects by users at a limited set of remote hosts. In this case the access control center provides the linkage between the connection oriented abstraction (as discussed in the Introduction) and the overall network security policy for DAC. In all cases the enforcement of the decision must be provided by the host where the object resides.

THERE ARE TWO FORMS OF DISTRIBUTION FOR THE DAC MECHANISM: IMPLEMENTING PORTIONS OF THE DAC IN SEPARATE COMPONENTS, AND SUPPORTING THE DAC IN SUBJECTS CONTAINED WITHIN THE NTCB PARTITION IN A COMPONENT. SINCE "THE ADP SYSTEM" IS UNDERSTOOD TO BE "THE COMPUTER NETWORK" AS A WHOLE, EACH NETWORK COMPONENT IS RESPONSIBLE FOR ENFORCING SECURITY IN THE MECHANISMS ALLOCATED TO IT TO ENSURE SECURE IMPLEMENTATION OF THE NETWORK SECURITY POLICY. FOR TRADITIONAL HOST SYSTEMS IT IS FREQUENTLY EASY TO ALSO ENFORCE THE DAC ALONG WITH THE MAC WITHIN THE REFERENCE MONITOR, PER SE, ALTHOUGH A FEW APPROACHES, SUCH AS VIRTUAL MACHINE MONITORS, SUPPORT DAC OUTSIDE THIS INTERFACE.

IN CONTRAST TO THE UNIVERSALLY RIGID STRUCTURE OF MANDATORY POLICIES (SEE THE MANDATORY ACCESS CONTROL SECTION), DAC POLICIES TEND TO BE VERY NETWORK AND SYSTEM SPECIFIC, WITH FEATURES THAT REFLECT THE NATURAL USE OF THE SYSTEM. FOR NETWORKS IT IS COMMON THAT INDIVIDUAL HOSTS WILL IMPOSE CONTROLS OVER THEIR LOCAL USERS ON THE BASIS OF NAMED INDIVIDUALS--MUCH LIKE THE CONTROLS USED WHEN THERE IS NO NETWORK CONNECTION. HOWEVER, IT IS DIFFICULT TO MANAGE IN A CENTRALIZED MANNER ALL THE INDIVIDUALS USING A LARGE NETWORK. THEREFORE, USERS ON OTHER HOSTS ARE COMMONLY GROUPED TOGETHER SO THAT THE CONTROLS REQUIRED BY THE NETWORK DAC POLICY ARE ACTUALLY BASED ON THE IDENTITY OF THE HOSTS OR OTHER COMPONENTS. A GATEWAY IS AN EXAMPLE OF SUCH A COMPONENT.

THE ASSURANCE REQUIREMENTS ARE AT THE VERY HEART OF THE CONCEPT OF A TRUSTED SYSTEM. IT IS THE ASSURANCE THAT DETERMINES IF A SYSTEM OR NETWORK IS APPROPRIATE FOR A GIVEN ENVIRONMENT, AS REFLECTED, FOR EXAMPLE, IN THE ENVIRONMENTS GUIDELINE-. IN THE CASE OF MONOLITHIC SYSTEMS THAT HAVE DAC INTEGRAL TO THE REFERENCE MONITOR, THE ASSURANCE REQUIREMENTS FOR DAC ARE INSEPARABLE FROM THOSE OF THE REST OF THE REFERENCE MONITOR. FOR NETWORKS THERE IS TYPICALLY A MUCH CLEARER DISTINCTION DUE TO DISTRIBUTED DAC. THE RATIONALE FOR MAKING THE DISTINCTION IN THIS NETWORK INTERPRETATION IS THAT IF MAJOR TRUSTED NETWORK COMPONENTS CAN BE MADE SIGNIFICANTLY EASIER TO DESIGN AND IMPLEMENT WITHOUT REDUCING THE ABILITY TO MEET SECURITY POLICY, THEN TRUSTED NETWORKS WILL BE MORE EASILY AVAILABLE.

3.1.1.2 Object Reuse

+ Statement from DoD 5200.28-STD

All authorizations to the information contained within a storage object shall be revoked prior to initial assignment, allocation or reallocation to a subject from the TCB's pool of unused storage objects. No information, including encrypted representations of information, produced by a prior subject's actions is to be available to any subject that obtains access to an object that has been released back to the system.

+ Interpretation

The NTCB shall ensure that any storage objects that it controls (e.g., message buffers under the control of a NTCB partition in a component) contain no information for which a subject in that component is not authorized before granting access. This requirement must be enforced by each of the NTCB partitions.

- Guidance for Applying the Department of Defense
Trusted Computer System Evaluation Criteria in Specific
Environments, CSC-STD-003-85.

+ Rationale

In a network system, storage objects of interest are things that the NTCB directly controls, such as message buffers in components. Each component of the network system

must enforce the object reuse requirement with respect to the storage objects of interest as determined by the network security policy. For example, the DAC requirement in this division leads to the requirement here that message buffers be under the control of the NTCB partition. A buffer assigned to an internal subject may be reused at the discretion of that subject which is responsible for preserving the integrity of message streams. Such controlled objects may be implemented in physical resources, such as buffers, disk sectors, tape space, and main memory, in components such as network switches.

3.1.1.3 Labels

+ Statement from DoD 5200.28-STD

SENSITIVITY LABELS ASSOCIATED WITH EACH SUBJECT AND STORAGE OBJECT UNDER ITS CONTROL (E.G., PROCESS, FILE, SEGMENT, DEVICE) SHALL BE MAINTAINED BY THE TCB. THESE LABELS SHALL BE USED AS THE BASIS FOR MANDATORY ACCESS CONTROL DECISIONS. IN ORDER TO IMPORT NON-LABELED DATA, THE TCB SHALL REQUEST AND RECEIVE FROM AN AUTHORIZED USER THE SENSITIVITY LEVEL OF THE DATA, AND ALL SUCH ACTIONS SHALL BE AUDITABLE BY THE TCB.

+ Interpretation

NON-LABELED DATA IMPORTED UNDER THE CONTROL OF THE NTCB PARTITION WILL BE ASSIGNED A LABEL CONSTRAINED BY THE SINGLE-LEVEL DEVICE USED TO IMPORT IT. LABELS MAY INCLUDE SECRECY AND INTEGRITY- COMPONENTS IN ACCORDANCE WITH THE OVERALL NETWORK SECURITY POLICY DESCRIBED BY THE NETWORK SPONSOR. WHENEVER THE TERM "LABEL" IS USED THROUGHOUT THIS INTERPRETATION, IT IS UNDERSTOOD TO INCLUDE BOTH COMPONENTS AS APPLICABLE. SIMILARLY, THE TERMS "SINGLE-LEVEL" AND "MULTILEVEL" ARE UNDERSTOOD TO BE BASED ON BOTH THE SECRECY AND INTEGRITY COMPONENTS OF THE POLICY. THE MANDATORY INTEGRITY POLICY WILL TYPICALLY HAVE REQUIREMENTS, SUCH AS THE PROBABILITY OF UNDETECTED MESSAGE STREAM MODIFICATION, THAT WILL BE REFLECTED IN THE LABEL FOR THE DATA SO PROTECTED. FOR EXAMPLE, WHEN DATA IS IMPORTED ITS INTEGRITY LABEL MAY BE ASSIGNED BASED ON MECHANISMS, SUCH AS CRYPTOGRAPHY, USED TO PROVIDE THE ASSURANCE REQUIRED BY THE POLICY. THE NTCB SHALL ASSURE THAT SUCH MECHANISM ARE PROTECTED FROM TAMPERING AND ARE ALWAYS INVOKED WHEN THEY ARE THE BASIS FOR

- See, for example, Biba, K.J., "Integrity Consideration for Secure Computer Systems," ESD-TR-76-372, MTR-3153, The MITRE Corporation, Bedford, MA, April 1977.

A LABEL.

+ Rationale

THE INTERPRETATION IS AN EXTENSION OF THE REQUIREMENT INTO THE CONTEXT OF A NETWORK SYSTEM AND PARTITIONED NTCB AS DEFINED FOR THESE NETWORK INTERPRETATIONS. A SINGLE-LEVEL DEVICE MAY BE REGARDED EITHER AS A SUBJECT OR AN OBJECT. A MULTILEVEL DEVICE IS REGARDED AS A TRUSTED SUBJECT IN WHICH THE SECURITY RANGE OF THE SUBJECT IS THE MINIMUM-MAXIMUM RANGE OF THE DATA EXPECTED TO BE TRANSMITTED OVER THE DEVICE.

THE SENSITIVITY LABELS FOR EITHER SECRECY OR INTEGRITY OR BOTH MAY REFLECT NON-HIERARCHICAL CATEGORIES OR HIERARCHICAL CLASSIFICATION OR BOTH.

3.1.1.3.1 Label Integrity

+ Statement from DoD 5200.28-STD

SENSITIVITY LABELS SHALL ACCURATELY REPRESENT SENSITIVITY LEVELS OF THE SPECIFIC SUBJECTS OR OBJECTS WITH WHICH THEY ARE ASSOCIATED. WHEN EXPORTED BY THE TCB, SENSITIVITY LABELS SHALL ACCURATELY AND UNAMBIGUOUSLY REPRESENT THE INTERNAL LABELS AND SHALL BE ASSOCIATED WITH THE INFORMATION BEING EXPORTED.

+ Interpretation

THE PHRASE "EXPORTED BY THE TCB" IS UNDERSTOOD TO INCLUDE TRANSMISSION OF INFORMATION FROM AN OBJECT IN ONE COMPONENT TO AN OBJECT IN ANOTHER COMPONENT. INFORMATION TRANSFERRED BETWEEN NTCB PARTITIONS IS ADDRESSED IN THE SYSTEM INTEGRITY SECTION. THE FORM OF INTERNAL AND EXTERNAL (EXPORTED) SENSITIVITY LABELS MAY DIFFER, BUT THE MEANING SHALL BE THE SAME. THE NTCB SHALL, IN ADDITION, ENSURE THAT CORRECT ASSOCIATION OF SENSITIVITY LABELS WITH THE INFORMATION BEING TRANSPORTED ACROSS THE NETWORK IS PRESERVED.

AS MENTIONED IN THE TRUSTED FACILITY MANUAL SECTION, ENCRYPTION TRANSFORMS THE REPRESENTATION OF INFORMATION SO THAT IT IS UNINTELLIGIBLE TO UNAUTHORIZED SUBJECTS. REFLECTING THIS TRANSFORMATION, THE SENSITIVITY LEVEL OF THE CIPHERTEXT IS GENERALLY LOWER THAN THE CLEARTEXT. IT FOLLOWS THAT CLEARTEXT AND CIPHERTEXT ARE CONTAINED IN DIFFERENT OBJECTS, EACH POSSESSING ITS OWN LABEL. THE LABEL OF THE CLEARTEXT MUST BE PRESERVED AND ASSOCIATED WITH THE CIPHERTEXT SO THAT IT CAN BE RESTORED WHEN THE CLEARTEXT IS SUBSEQUENTLY OBTAINED BY DECRYPTING THE CIPHERTEXT. IF THE CLEARTEXT IS ASSOCIATED WITH A SINGLE-LEVEL DEVICE, THE LABEL OF THAT CLEARTEXT MAY BE IMPLICIT. THE LABEL MAY ALSO BE IMPLICIT IN THE KEY.

WHEN INFORMATION IS EXPORTED TO AN ENVIRONMENT WHERE IT IS SUBJECT TO DELIBERATE OR ACCIDENTAL MODIFICATION, THE TCB SHALL SUPPORT THE MEANS, SUCH AS CRYPTOGRAPHIC CHECKSUMS, TO ASSURE THE ACCURACY OF THE LABELS. WHEN THERE IS A MANDATORY INTEGRITY POLICY, THE POLICY WILL DEFINE THE MEANING OF INTEGRITY LABELS.

+ Rationale

ENCRYPTION ALGORITHMS AND THEIR IMPLEMENTATION ARE OUTSIDE THE SCOPE OF THESE INTERPRETATIONS. SUCH ALGORITHMS MAY BE IMPLEMENTED IN A SEPARATE DEVICE OR MAY BE INCORPORATED IN A SUBJECT OF A LARGER COMPONENT. WITHOUT PREJUDICE, EITHER IMPLEMENTATION PACKAGING IS REFERRED TO AS AN ENCRYPTION MECHANISM HEREIN. IF ENCRYPTION METHODOLOGIES ARE EMPLOYED IN THIS REGARD, THEY SHALL BE APPROVED BY THE NATIONAL SECURITY AGENCY (NSA). THE ENCRYPTION PROCESS IS PART OF THE NETWORK TRUSTED COMPUTER BASE PARTITION IN THE COMPONENTS IN WHICH IT IS IMPLEMENTED.

THE ENCRYPTION MECHANISM IS NOT NECESSARILY A MULTILEVEL DEVICE OR MULTILEVEL SUBJECT, AS THESE TERMS ARE USED IN THESE CRITERIA. THE PROCESS OF ENCRYPTION IS MULTILEVEL BY DEFINITION. THE CLEARTEXT AND CIPHERTEXT INTERFACES CARRY INFORMATION OF DIFFERENT SENSITIVITY. AN ENCRYPTION MECHANISM DOES NOT PROCESS DATA IN THE SENSE OF PERFORMING LOGICAL OR ARITHMETIC OPERATIONS ON THAT DATA WITH THE INTENT OF PRODUCING NEW DATA. THE CLEARTEXT AND CIPHERTEXT INTERFACES ON THE ENCRYPTION MECHANISM MUST BE SEPARATELY IDENTIFIED AS BEING SINGLE-LEVEL OR MULTILEVEL. IF THE INTERFACE IS SINGLE-LEVEL, THEN THE SENSITIVITY OF THE DATA IS ESTABLISHED BY A TRUSTED INDIVIDUAL AND IMPLICITLY ASSOCIATED WITH THE INTERFACE; THE EXPORTATION TO SINGLE-LEVEL DEVICES CRITERION APPLIES.

IF THE INTERFACE IS MULTILEVEL, THEN THE DATA MUST BE LABELED; THE EXPORTATION TO MULTILEVEL DEVICES CRITERION APPLIES. THE NETWORK ARCHITECT IS FREE TO SELECT AN ACCEPTABLE MECHANISM FOR ASSOCIATING A LABEL WITH AN OBJECT. WITH REFERENCE TO ENCRYPTED OBJECTS, THE FOLLOWING EXAMPLES ARE POSSIBLE:

1. INCLUDE A LABEL FIELD IN THE PROTOCOL DEFINITION OF THE OBJECT.
2. IMPLICITLY ASSOCIATE THE LABEL WITH THE OBJECT THROUGH THE ENCRYPTION KEY. THAT IS, THE ENCRYPTION KEY UNIQUELY IDENTIFIES A SENSITIVITY LEVEL. A SINGLE OR PRIVATE KEY MUST BE PROTECTED AT THE LEVEL OF THE DATA THAT IT ENCRYPTS.

3.1.1.3.2 Exportation of Labeled Information

+ Statement from DoD 5200.28-STD

THE TCB SHALL DESIGNATE EACH COMMUNICATION CHANNEL AND I/O DEVICE AS EITHER SINGLE-LEVEL OR MULTILEVEL. ANY CHANGE IN THIS DESIGNATION SHALL BE DONE MANUALLY AND SHALL BE AUDITABLE BY THE TCB. THE TCB SHALL MAINTAIN AND BE ABLE TO AUDIT ANY CHANGE IN THE SENSITIVITY LEVEL OR LEVELS ASSOCIATED WITH A COMMUNICATIONS CHANNEL OR I/O DEVICE.

+ Interpretation

EACH COMMUNICATION CHANNEL AND NETWORK COMPONENT SHALL BE DESIGNATED AS EITHER SINGLE-LEVEL OR MULTILEVEL. ANY CHANGE IN THIS DESIGNATION SHALL BE DONE WITH THE COGNIZANCE AND APPROVAL OF THE ADMINISTRATOR OR SECURITY OFFICER IN CHARGE OF THE AFFECTED COMPONENTS AND THE ADMINISTRATOR OR SECURITY OFFICER IN CHARGE OF THE NTCB. THIS CHANGE SHALL BE AUDITABLE BY THE NETWORK. THE NTCB SHALL MAINTAIN AND BE ABLE TO AUDIT ANY CHANGE IN THE CURRENT SENSITIVITY LEVEL ASSOCIATED WITH THE DEVICE CONNECTED TO A SINGLE-LEVEL COMMUNICATION CHANNEL OR THE RANGE ASSOCIATED WITH A MULTILEVEL COMMUNICATION CHANNEL OR COMPONENT. THE NTCB SHALL ALSO BE ABLE TO AUDIT ANY CHANGE IN THE SET OF SENSITIVITY LEVELS ASSOCIATED WITH THE INFORMATION WHICH CAN BE TRANSMITTED OVER A MULTILEVEL COMMUNICATION CHANNEL OR COMPONENT.

+ Rationale

COMMUNICATION CHANNELS AND COMPONENTS IN A NETWORK ARE ANALOGOUS TO COMMUNICATION CHANNELS AND I/O DEVICES IN STAND-ALONE SYSTEMS. THEY MUST BE DESIGNATED AS EITHER MULTILEVEL (I.E., ABLE TO DISTINGUISH AND MAINTAIN SEPARATION AMONG INFORMATION OF VARIOUS SENSITIVITY LEVELS) OR SINGLE-LEVEL. AS IN THE TCSEC, SINGLE-LEVEL DEVICES MAY ONLY BE ATTACHED TO SINGLE-LEVEL CHANNELS.

THE LEVEL OR SET OF LEVELS OF INFORMATION THAT CAN BE SENT TO A COMPONENT OR OVER A COMMUNICATION CHANNEL SHALL ONLY CHANGE WITH THE KNOWLEDGE AND APPROVAL OF THE SECURITY OFFICERS (OR SYSTEM ADMINISTRATOR, IF THERE IS NO SECURITY OFFICER) OF THE NETWORK, AND OF THE AFFECTED COMPONENTS. THIS REQUIREMENT ENSURES THAT NO SIGNIFICANT SECURITY-RELEVANT CHANGES ARE MADE WITHOUT THE APPROVAL OF ALL AFFECTED PARTIES.

3.1.1.3.2.1 Exportation to Multilevel Devices

+ Statement from DoD 5200.28-STD

WHEN THE TCB EXPORTS AN OBJECT TO A MULTILEVEL I/O DEVICE, THE SENSITIVITY LABEL ASSOCIATED WITH THAT OBJECT SHALL ALSO BE EXPORTED AND SHALL RESIDE ON THE SAME PHYSICAL MEDIUM AS

THE EXPORTED INFORMATION AND SHALL BE IN THE SAME FORM (I.E., MACHINE-READABLE OR HUMAN-READABLE FORM). WHEN THE TCB EXPORTS OR IMPORTS AN OBJECT OVER A MULTILEVEL COMMUNICATIONS CHANNEL, THE PROTOCOL USED ON THAT CHANNEL SHALL PROVIDE FOR THE UNAMBIGUOUS PAIRING BETWEEN THE SENSITIVITY LABELS AND THE ASSOCIATED INFORMATION THAT IS SENT OR RECEIVED.

+ Interpretation

THE COMPONENTS, INCLUDING HOSTS, OF A NETWORK SHALL BE INTERCONNECTED OVER "MULTILEVEL COMMUNICATION CHANNELS," MULTIPLE SINGLE-LEVEL COMMUNICATION CHANNELS, OR BOTH, WHENEVER THE INFORMATION IS TO BE PROTECTED AT MORE THAN A SINGLE SENSITIVITY LEVEL. THE PROTOCOL FOR ASSOCIATING THE SENSITIVITY LABEL AND THE EXPORTED INFORMATION SHALL PROVIDE THE ONLY INFORMATION NEEDED TO CORRECTLY ASSOCIATE A SENSITIVITY LEVEL WITH THE EXPORTED INFORMATION TRANSFERRED OVER THE MULTILEVEL CHANNEL BETWEEN THE NTCB PARTITIONS IN INDIVIDUAL COMPONENTS. THIS PROTOCOL DEFINITION MUST SPECIFY THE REPRESENTATION AND SEMANTICS OF THE SENSITIVITY LABELS (I.E., THE MACHINE-READABLE LABEL MUST UNIQUELY REPRESENT THE SENSITIVITY LEVEL).

THE "UNAMBIGUOUS" ASSOCIATION OF THE SENSITIVITY LEVEL WITH THE COMMUNICATED INFORMATION SHALL MEET THE SAME LEVEL OF ACCURACY AS THAT REQUIRED FOR ANY OTHER LABEL WITHIN THE NTCB, AS SPECIFIED IN THE CRITERION FOR LABEL INTEGRITY. THIS MAY BE PROVIDED BY PROTECTED AND HIGHLY RELIABLE DIRECT PHYSICAL LAYER CONNECTIONS, OR BY TRADITIONAL CRYPTOGRAPHIC LINK PROTECTION IN WHICH ANY ERRORS DURING TRANSMISSION CAN BE READILY DETECTED, OR BY USE OF A SEPARATE CHANNEL.

+ Rationale

THIS PROTOCOL MUST SPECIFY THE REPRESENTATION AND SEMANTICS OF THE SENSITIVITY LABELS. SEE THE MANDATORY ACCESS CONTROL POLICIES SECTION IN APPENDIX B. THE MULTILEVEL DEVICE INTERFACE TO (UNTRUSTED) SUBJECTS MAY BE IMPLEMENTED EITHER BY THE INTERFACE OF THE REFERENCE MONITOR, PER SE, OR BY A MULTILEVEL SUBJECT (E.G., A "TRUSTED SUBJECT" AS DEFINED IN THE BELL-LAPADULA MODEL) THAT PROVIDES THE LABELS BASED ON THE INTERNAL LABELS OF THE NTCB PARTITION.

THE CURRENT STATE OF THE ART LIMITS THE SUPPORT FOR MANDATORY POLICY THAT IS PRACTICAL FOR SECURE NETWORKS. REFERENCE MONITOR SUPPORT TO ENSURE THE CONTROL OVER ALL THE OPERATIONS OF EACH SUBJECT IN THE NETWORK MUST BE COMPLETELY PROVIDED WITHIN THE SINGLE NTCB PARTITION ON WHICH THAT SUBJECT INTERFACES TO THE NTCB. THIS MEANS THAT THE ENTIRE PORTION OF THE "SECURE STATE" REPRESENTED IN THE SECURITY POLICY MODEL THAT MAY BE CHANGED BY TRANSITIONS INVOKED BY THIS SUBJECT MUST BE CONTAINED IN THE SAME COMPONENT.

THE SECURE STATE OF AN NTCB PARTITION MAY BE AFFECTED BY EVENTS EXTERNAL TO THE COMPONENT IN WHICH THE NTCB PARTITION RESIDES (E.G., ARRIVAL OF A MESSAGE). THE EFFECT OCCURS ASYNCHRONOUSLY AFTER BEING INITIATED BY AN EVENT IN ANOTHER COMPONENT OR PARTITION. FOR EXAMPLE, INDETERMINATE DELAYS MAY OCCUR BETWEEN THE INITIATION OF A MESSAGE IN ONE COMPONENT, THE ARRIVAL OF THE MESSAGE IN THE NTCB PARTITION IN ANOTHER COMPONENT, AND THE CORRESPONDING CHANGE TO THE SECURE STATE OF THE SECOND COMPONENT. SINCE EACH COMPONENT IS EXECUTING CONCURRENTLY, TO DO OTHERWISE WOULD REQUIRE SOME SORT OF NETWORK-WIDE CONTROL TO SYNCHRONIZE STATE TRANSITIONS, SUCH AS A GLOBAL NETWORK-WIDE CLOCK FOR ALL PROCESSORS; IN GENERAL, SUCH DESIGNS ARE NOT PRACTICAL AND PROBABLY NOT EVEN DESIRABLE. THEREFORE, THE INTERACTION BETWEEN NTCB PARTITIONS IS RESTRICTED TO JUST COMMUNICATIONS BETWEEN PAIRS (AT LEAST LOGICALLY) OF DEVICES-MULTILEVEL DEVICES IF THE DEVICE(S) CAN SEND/RECEIVE DATA OF MORE THAN A SINGLE LEVEL. FOR BROADCAST CHANNELS THE PAIRS ARE THE SENDER AND INTENDED RECEIVER(S). HOWEVER, IF THE BROADCAST CHANNEL CARRIES MULTIPLE LEVELS OF INFORMATION, ADDITIONAL MECHANISM (E.G., CRYPTOCHECKSUM MAINTAINED BY THE TCB) MAY BE REQUIRED TO ENFORCE SEPARATION AND PROPER DELIVERY.

A COMMON REPRESENTATION FOR SENSITIVITY LABELS IS NEEDED IN THE PROTOCOL USED ON THAT CHANNEL AND UNDERSTOOD BY BOTH THE SENDER AND RECEIVER WHEN TWO MULTILEVEL DEVICES (IN THIS CASE, IN TWO DIFFERENT COMPONENTS) ARE INTERCONNECTED. EACH DISTINCT SENSITIVITY LEVEL OF THE OVERALL NETWORK POLICY MUST BE REPRESENTED UNIQUELY IN THESE LABELS.

WITHIN A MONOLITHIC TCB, THE ACCURACY OF THE SENSITIVITY LABELS IS GENERALLY ASSURED BY SIMPLE TECHNIQUES, E.G., VERY RELIABLE CONNECTIONS OVER VERY SHORT PHYSICAL CONNECTIONS, SUCH AS ON A SINGLE PRINTED CIRCUIT BOARD OR OVER AN INTERNAL BUS. IN MANY NETWORK ENVIRONMENTS THERE IS A MUCH HIGHER PROBABILITY OF ACCIDENTALLY OR MALICIOUSLY INTRODUCED ERRORS, AND THESE MUST BE PROTECTED AGAINST.

3.1.1.3.2.2 Exportation to Single-Level Devices

+ Statement from DoD 5200.28-STD

SINGLE-LEVEL I/O DEVICES AND SINGLE-LEVEL COMMUNICATION CHANNELS ARE NOT REQUIRED TO MAINTAIN THE SENSITIVITY LABELS OF THE INFORMATION THEY PROCESS. HOWEVER, THE TCB SHALL INCLUDE A MECHANISM BY WHICH THE TCB AND AN AUTHORIZED USER RELIABLY COMMUNICATE TO DESIGNATE THE SINGLE SENSITIVITY LEVEL OF INFORMATION IMPORTED OR EXPORTED VIA SINGLE-LEVEL COMMUNICATION CHANNELS OR I/O DEVICES.

+ Interpretation

WHENEVER ONE OR BOTH OF TWO DIRECTLY CONNECTED COMPONENTS IS NOT TRUSTED TO MAINTAIN THE SEPARATION OF INFORMATION OF DIFFERENT SENSITIVITY LEVELS, OR WHENEVER THE TWO DIRECTLY CONNECTED COMPONENTS HAVE ONLY A SINGLE SENSITIVITY LEVEL IN COMMON, THE TWO COMPONENTS OF THE NETWORK SHALL COMMUNICATE OVER A SINGLE-LEVEL CHANNEL. SINGLE-LEVEL COMPONENTS AND SINGLE-LEVEL COMMUNICATION CHANNELS ARE NOT REQUIRED TO MAINTAIN THE SENSITIVITY LABELS OF THE INFORMATION THEY PROCESS. HOWEVER, THE NTCB SHALL INCLUDE A RELIABLE COMMUNICATION MECHANISM BY WHICH THE NTCB AND AN AUTHORIZED USER OR A SUBJECT WITHIN AN NTCB PARTITION CAN DESIGNATE THE SINGLE SENSITIVITY LEVEL OF INFORMATION IMPORTED OR EXPORTED VIA SINGLE-LEVEL COMMUNICATION CHANNELS OR NETWORK COMPONENTS.

+ Rationale

SINGLE-LEVEL COMMUNICATIONS CHANNELS AND SINGLE-LEVEL COMPONENTS IN NETWORKS ARE ANALOGOUS TO SINGLE LEVEL CHANNELS AND I/O DEVICES IN STAND-ALONE SYSTEMS IN THAT THEY ARE NOT TRUSTED TO MAINTAIN THE SEPARATION OF INFORMATION OF DIFFERENT SENSITIVITY LEVELS. THE LABELS ASSOCIATED WITH DATA TRANSMITTED OVER THOSE CHANNELS AND BY THOSE COMPONENTS ARE THEREFORE IMPLICIT; THE NTCB ASSOCIATES LABELS WITH THE DATA BECAUSE OF THE CHANNEL OR COMPONENT, NOT BECAUSE OF AN EXPLICIT PART OF THE BIT STREAM. NOTE THAT THE SENSITIVITY LEVEL OF ENCRYPTED INFORMATION IS THE LEVEL OF THE CIPHER-TEXT RATHER THAN THE ORIGINAL LEVEL(S) OF THE PLAINTEXT.

3.1.1.3.2.3 Labeling Human-Readable Output

+ Statement from DoD 5200.28-STD

THE ADP SYSTEM ADMINISTRATOR SHALL BE ABLE TO SPECIFY THE PRINTABLE LABEL NAMES ASSOCIATED WITH EXPORTED SENSITIVITY LABELS. THE TCB SHALL MARK THE BEGINNING AND END OF ALL HUMAN-READABLE, PAGED, HARDCOPY OUTPUT (E.G., LINE PRINTER OUTPUT) WITH HUMAN-READABLE SENSITIVITY LABELS THAT PROPERLY REPRESENT THE SENSITIVITY OF THE OUTPUT. THE TCB SHALL, BY DEFAULT, MARK THE TOP AND BOTTOM OF EACH PAGE OF HUMAN-READABLE, PAGED, HARDCOPY OUTPUT (E.G., LINE PRINTER OUTPUT) WITH HUMAN-READABLE SENSITIVITY LABELS THAT PROPERLY REPRESENT THE SENSITIVITY OF THE PAGE. THE TCB SHALL, BY DEFAULT AND IN AN APPROPRIATE MANNER, MARK OTHER FORMS OF HUMAN READABLE OUTPUT (E.G., MAPS, GRAPHICS) WITH HUMAN-READABLE SENSITIVITY LABELS THAT PROPERLY REPRESENT THE SENSITIVITY OF THE OUTPUT. ANY OVERRIDE OF THESE MARKINGS DEFAULTS SHALL BE AUDITABLE BY THE TCB.

+ Interpretation

THIS CRITERION IMPOSES NO REQUIREMENT TO A COMPONENT THAT PRODUCES NO HUMAN-READABLE OUTPUT. FOR THOSE THAT DO

PRODUCE HUMAN-READABLE OUTPUT, EACH SENSITIVITY LEVEL THAT

1 THE HIERARCHICAL CLASSIFICATION COMPONENT IN HUMAN-READABLE SENSITIVITY LABELS SHALL BE EQUAL TO THE GREATEST HIERARCHICAL CLASSIFICATION OF ANY OF THE INFORMATION IN THE OUTPUT THAT THE LABELS REFER TO; THE NON-HIERARCHICAL CATEGORY COMPONENT SHALL INCLUDE ALL OF THE NON-HIERARCHICAL CATEGORIES OF THE INFORMATION IN THE OUTPUT THE LABELS REFER TO, BUT NO OTHER NON-HIERARCHICAL CATEGORIES.

IS DEFINED TO THE NETWORK SHALL HAVE A UNIFORM MEANING ACROSS ALL COMPONENTS. THE NETWORK ADMINISTRATOR, IN CONJUNCTION WITH ANY AFFECTED COMPONENT ADMINISTRATOR, SHALL BE ABLE TO SPECIFY THE HUMAN-READABLE LABEL THAT IS ASSOCIATED WITH EACH DEFINED SENSITIVITY LEVEL.

+ Rationale

THE INTERPRETATION IS A STRAIGHTFORWARD EXTENSION OF THE REQUIREMENT INTO THE CONTEXT OF A NETWORK SYSTEM AND PARTITIONED NTCB AS DEFINED FOR THESE NETWORK INTERPRETATIONS.

3.1.1.4 Mandatory Access Control

+ Statement from DoD 5200.28-STD

THE TCB SHALL ENFORCE A MANDATORY ACCESS CONTROL POLICY OVER ALL SUBJECTS AND STORAGE OBJECTS UNDER ITS CONTROL (E.G., PROCESSES, FILES, SEGMENTS, DEVICES). THESE SUBJECTS AND OBJECTS SHALL BE ASSIGNED SENSITIVITY LABELS THAT ARE A COMBINATION OF HIERARCHICAL CLASSIFICATION LEVELS AND NON-HIERARCHICAL CATEGORIES, AND THE LABELS SHALL BE USED AS THE BASIS FOR MANDATORY ACCESS CONTROL DECISIONS. THE TCB SHALL BE ABLE TO SUPPORT TWO OR MORE SUCH SENSITIVITY LEVELS. (SEE THE MANDATORY ACCESS CONTROL INTERPRETATIONS.) THE FOLLOWING REQUIREMENTS SHALL HOLD FOR ALL ACCESSES BETWEEN SUBJECTS AND OBJECTS CONTROLLED BY THE TCB: A SUBJECT CAN READ AN OBJECT ONLY IF THE HIERARCHICAL CLASSIFICATION IN THE SUBJECT'S SENSITIVITY LEVEL IS GREATER THAN OR EQUAL TO THE HIERARCHICAL CLASSIFICATION OF THE OBJECT'S SENSITIVITY LEVEL AND THE NON-HIERARCHICAL CATEGORIES IN THE SUBJECT'S SENSITIVITY LEVEL INCLUDE ALL THE NON-HIERARCHICAL CATEGORIES IN THE OBJECT'S SENSITIVITY LEVEL. A SUBJECT CAN WRITE AN OBJECT ONLY IF THE HIERARCHICAL CLASSIFICATION IN THE SUBJECT'S SENSITIVITY LEVEL IS LESS THAN OR EQUAL TO THE HIERARCHICAL CLASSIFICATION OF THE OBJECT'S SENSITIVITY LEVEL AND THE NON-HIERARCHICAL CATEGORIES IN THE SUBJECT'S

SENSITIVITY LEVEL ARE INCLUDED IN THE NON-HIERARCHICAL CATEGORIES IN THE OBJECT'S SENSITIVITY LEVEL. IDENTIFICATION AND AUTHENTICATION DATA SHALL BE USED BY THE TCB TO AUTHENTICATE THE USER'S IDENTITY AND TO ENSURE THAT THE SENSITIVITY LEVEL AND AUTHORIZATION OF SUBJECTS EXTERNAL TO THE TCB THAT MAY BE CREATED TO ACT ON BEHALF OF THE INDIVIDUAL USER ARE DOMINATED BY THE CLEARANCE AND AUTHORIZATION OF THAT USER.

+ Interpretation

EACH PARTITION OF THE NTCB EXERCISES MANDATORY ACCESS CONTROL POLICY OVER ALL SUBJECTS AND OBJECTS IN ITS COMPONENT THAT ARE UNDER ITS CONTROL. IN A NETWORK, THE RESPONSIBILITY OF AN NTCB PARTITION ENCOMPASSES ALL MANDATORY ACCESS CONTROL FUNCTIONS IN ITS COMPONENT THAT WOULD BE REQUIRED OF A TCB IN A STAND-ALONE SYSTEM. IN PARTICULAR, SUBJECTS AND OBJECTS USED FOR COMMUNICATION WITH OTHER COMPONENTS ARE UNDER THE CONTROL OF THE NTCB PARTITION. MANDATORY ACCESS CONTROL INCLUDES SECRECY AND INTEGRITY CONTROL TO THE EXTENT THAT THE NETWORK SPONSOR HAS DESCRIBED IN THE OVERALL NETWORK SECURITY POLICY.

CONCEPTUAL ENTITIES ASSOCIATED WITH COMMUNICATION BETWEEN TWO COMPONENTS, SUCH AS SESSIONS, CONNECTIONS AND VIRTUAL CIRCUITS, MAY BE THOUGHT OF AS HAVING TWO ENDS, ONE IN EACH COMPONENT, WHERE EACH END IS REPRESENTED BY A LOCAL OBJECT. COMMUNICATION IS VIEWED AS AN OPERATION THAT COPIES INFORMATION FROM AN OBJECT AT ONE END OF A COMMUNICATION PATH TO AN OBJECT AT THE OTHER END. TRANSIENT DATA-CARRYING ENTITIES, SUCH AS DATAGRAMS AND PACKETS, EXIST EITHER AS INFORMATION WITHIN OTHER OBJECTS, OR AS A PAIR OF OBJECTS, ONE AT EACH END OF THE COMMUNICATION PATH.

THE REQUIREMENT FOR "TWO OR MORE" SENSITIVITY LEVELS CAN BE MET BY EITHER SECRECY OR INTEGRITY LEVELS. WHEN THERE IS A MANDATORY INTEGRITY POLICY, THE STATED REQUIREMENTS FOR READING AND WRITING ARE GENERALIZED TO: A SUBJECT CAN READ AN OBJECT ONLY IF THE SUBJECT'S SENSITIVITY LEVEL DOMINATES THE OBJECT'S SENSITIVITY LEVEL, AND A SUBJECT CAN WRITE AN OBJECT ONLY IF THE OBJECT'S SENSITIVITY LEVEL DOMINATES THE SUBJECT'S SENSITIVITY LEVEL. BASED ON THE INTEGRITY POLICY, THE NETWORK SPONSOR SHALL DEFINE THE DOMINANCE RELATION FOR THE TOTAL LABEL, FOR EXAMPLE, BY COMBINING SECRECY AND INTEGRITY LATTICES. -

+ Rationale

AN NTCB PARTITION CAN MAINTAIN ACCESS CONTROL ONLY OVER SUBJECTS AND OBJECTS IN ITS COMPONENT. ACCESS BY A SUBJECT IN ONE COMPONENT TO INFORMATION CONTAINED IN AN OBJECT IN ANOTHER COMPONENT REQUIRES THE CREATION OF A SUBJECT IN THE REMOTE COMPONENT WHICH ACTS AS A SURROGATE FOR THE FIRST SUBJECT.

THE MANDATORY ACCESS CONTROLS MUST BE ENFORCED AT THE INTERFACE OF THE REFERENCE MONITOR (VIZ. THE MECHANISM THAT CONTROLS PHYSICAL PROCESSING RESOURCES) FOR EACH NTCB PARTITION. THIS MECHANISM CREATES THE ABSTRACTION OF SUBJECTS AND OBJECTS WHICH IT CONTROLS. SOME OF THESE SUBJECTS OUTSIDE THE REFERENCE MONITOR, PER SE, MAY BE DESIGNATED TO IMPLEMENT PART OF AN NTCB PARTITION'S MANDATORY POLICY, E.G., BY USING THE "TRUSTED SUBJECTS" DEFINED IN THE BELL-

- See, for example, Grohn, M. J., A Model of a Protected Data Management System, ESD-TR-76-289, I. P. Sharp Assoc. Ltd., June, 1976; and Denning, D. E., Lunt, T. F., Neumann, P. G., Schell, R. R., Heckman, M. and Shockley, W., Secure Distributed Data Views, Security Policy and Interpretation for a Class A1 Multilevel Secure Relational Database System, SRI International, November 1986.

LAPADULA MODEL.

THE PRIOR REQUIREMENTS ON EXPORTATION OF LABELED INFORMATION TO AND FROM I/O DEVICES ENSURE THE CONSISTENCY BETWEEN THE SENSITIVITY LABELS OF OBJECTS CONNECTED BY A COMMUNICATION PATH. AS NOTED IN THE INTRODUCTION, THE NETWORK ARCHITECTURE MUST RECOGNIZE THE LINKAGE BETWEEN THE OVERALL MANDATORY NETWORK SECURITY POLICY AND THE CONNECTION ORIENTED ABSTRACTION. FOR EXAMPLE, INDIVIDUAL DATA-CARRYING ENTITIES SUCH AS DATAGRAMS CAN HAVE INDIVIDUAL SENSITIVITY LABELS THAT SUBJECT THEM TO MANDATORY ACCESS CONTROL IN EACH COMPONENT. THE ABSTRACTION OF A SINGLE-LEVEL CONNECTION IS REALIZED AND ENFORCED IMPLICITLY BY AN ARCHITECTURE WHILE A CONNECTION IS REALIZED BY SINGLE-LEVEL SUBJECTS THAT NECESSARILY EMPLOY ONLY DATAGRAMS OF THE SAME LEVEL.

THE FUNDAMENTAL TRUSTED SYSTEMS TECHNOLOGY PERMITS THE DAC MECHANISM TO BE DISTRIBUTED, IN CONTRAST TO THE REQUIREMENTS FOR MANDATORY ACCESS CONTROL. FOR NETWORKS THIS SEPARATION OF MAC AND DAC MECHANISMS IS THE RULE RATHER THAN THE EXCEPTION.

THE SET OF TOTAL SENSITIVITY LABELS USED TO REPRESENT ALL THE SENSITIVITY LEVELS FOR THE MANDATORY ACCESS CONTROL (COMBINED DATA SECRECY AND DATA INTEGRITY) POLICY ALWAYS FORMS A PARTIALLY ORDERED SET. WITHOUT LOSS OF GENERALITY, THIS SET OF LABELS CAN ALWAYS BE EXTENDED TO FORM A LATTICE, BY INCLUDING ALL THE COMBINATIONS OF NON-HIERARCHICAL CATEGORIES. AS FOR ANY LATTICE, A DOMINANCE RELATION IS

ALWAYS DEFINED FOR THE TOTAL SENSITIVITY LABELS. FOR ADMINISTRATIVE REASONS IT MAY BE HELPFUL TO HAVE A MAXIMUM LEVEL WHICH DOMINATES ALL OTHERS.

3.1.2 Accountability

3.1.2.1 Identification and Authentication

+ Statement from DoD 5200.28-STD

The TCB shall require users to identify themselves to it before beginning to perform any other actions that the TCB is expected to mediate. Furthermore, the TCB shall MAINTAIN AUTHENTICATION DATA THAT INCLUDES INFORMATION FOR VERIFYING THE IDENTIFY OF INDIVIDUAL USERS (E.G., PASSWORDS) AS WELL AS INFORMATION FOR DETERMINING THE CLEARANCE AND AUTHORIZATIONS OF INDIVIDUAL USERS. THIS DATA SHALL BE USED BY THE TCB TO AUTHENTICATE THE USER'S IDENTITY AND TO ENSURE THAT THE SENSITIVITY LEVEL AND AUTHORIZATION OF SUBJECTS EXTERNAL TO THE TCB THAT MAY BE CREATED TO ACT ON BEHALF OF THE INDIVIDUAL USER ARE DOMINATED BY THE CLEARANCE AND AUTHORIZATION OF THAT USER. The TCB shall protect authentication data so that it cannot be accessed by any unauthorized user. The TCB shall be able to enforce individual accountability by providing the capability to uniquely identify each individual ADP system user. The TCB shall also provide the capability of associating this identity with all auditable actions taken by that individual.

+ Interpretation

The requirement for identification and authentication of users is the same for a network system as for an ADP system. The identification and authentication may be done by the component to which the user is directly connected or some other component, such as an identification and authentication server. Available techniques, such as those described in the Password Guideline, are generally also applicable in the network context. However, in cases where the NTCB is expected to mediate actions of a host (or other network component) that is acting on behalf of a user or group of users, the NTCB may employ identification and authentication of the host (or other component) in lieu of identification and authentication of an individual user, so long as the component identifier implies a list of specific users uniquely associated with the identifier at the time of its use for authentication. This requirement does not apply to internal subjects.

Authentication information, including the identity of a user (once authenticated) may be passed from one component

to another without reauthentication, so long as the NTCB protects (e.g., by encryption) the information from unauthorized disclosure and modification. This protection shall provide at least a similar level of assurance (or strength of mechanism) as pertains to the protection of the authentication mechanism and authentication data.

+ Rationale

The need for accountability is not changed in the context of a network system. The fact that the NTCB is partitioned over a set of components neither reduces the need nor imposes new requirements. That is, individual accountability is still the objective. Also, in the context of a network system at the (C2) level or higher "individual accountability" can be satisfied by identification of a host (or other component) so long as the requirement for traceability to individual users or a set of specific individual users with active subjects is satisfied. There is allowed to be an uncertainty in traceability because of elapsed time between changes in the group membership and the enforcement in the access control mechanisms. In addition, there is no need in a distributed processing system like a network to reauthenticate a user at each point in the network where a projection of a user (via the subject operating on behalf of the user) into another remote subject takes place.

= Department of Defense Password Management Guide-
line, CSC-STD-002-85

The passing of identifiers and/or authentication information from one component to another is usually done in support to the implementation of the discretionary access control (DAC). This support relates directly to the DAC regarding access by a user to a storage object in a different NTCB partition than the one where the user was authenticated. Employing a forwarded identification implies additional reliance on the source and components along the path. IF THE AUTHENTICATED IDENTIFICATION IS USED AS THE BASIS OF DETERMINING A SENSITIVITY LABEL FOR A SUBJECT, IT MUST SATISFY THE LABEL INTEGRITY CRITERION.

AN AUTHENTICATED IDENTIFICATION MAY BE FORWARDED BETWEEN COMPONENTS AND EMPLOYED IN SOME COMPONENT TO IDENTIFY THE SENSITIVITY LEVEL ASSOCIATED WITH A SUBJECT CREATED TO ACT ON BEHALF OF THE USER SO IDENTIFIED.

3.1.2.2 Audit

+ Statement from DoD 5200.28-STD

The TCB shall be able to create, maintain, and protect from modification or unauthorized access or destruction an audit trail of accesses to the objects it protects. The audit data shall be protected by the TCB so that read access to it is limited to those who are authorized for audit data. The TCB shall be able to record the following types of events: use of identification and authentication mechanisms, introduction of objects into a user's address space (e.g., file open, program initiation), deletion of objects, actions taken by computer operators and system administrators and/or system security officers, and other security relevant events. THE TCB SHALL ALSO BE ABLE TO AUDIT ANY OVERRIDE OF HUMAN-READABLE OUTPUT MARKINGS. For each recorded event, the audit record shall identify: date and time of the event, user, type of event, and success or failure of the event. For identification/authentication events the origin of request (e.g., terminal ID) shall be included in the audit record. For events that introduce an object into a user's address space and for object deletion events the audit record shall include the name of the object AND THE OBJECT'S SENSITIVITY LEVEL. The ADP system administrator shall be able to selectively audit the actions of any one or more users based on individual identify AND/OR OBJECT SENSITIVITY LEVEL.

+ Interpretation

This criterion applies as stated. The sponsor must select which events are auditable. If any such events are not distinguishable by the NTCB alone (for example those identified in Part II), the audit mechanism shall provide an interface, which an authorized subject can invoke with parameters sufficient to produce an audit record. These audit records shall be distinguishable from those provided by the NTCB. In the context of a network system, "other security relevant events" (depending on network system architecture and network security policy) might be as follows:

1. Identification of each access event (e.g., establishing a connection or a connectionless association between processes in two hosts of the network) and its principal parameters (e.g., host identifiers of the two hosts involved in the access event and user identifier or host identifier of the user or host that is requesting the access event)
2. Identification of the starting and ending times of each access event using local time or global syn-

chronized time

3. Identification of security-relevant exceptional conditions (e.g., potential violation of data integrity, such as misrouted datagrams) detected during the transactions between two hosts
4. Utilization of cryptographic variables
5. Changing the configuration of the network (e.g., a component leaving the network and rejoining)

In addition, identification information should be included in appropriate audit trail records, as necessary, to allow association of all related (e.g., involving the same network event) audit trail records (e.g., at different hosts) with each other. Furthermore, a component of the network system may provide the required audit capability (e.g., storage, retrieval, reduction, analysis) for other components that do not internally store audit data but transmit the audit data to some designated collection component. Provisions shall be made to control the loss of audit data due to unavailability of resources.

In the context of a network system, the "user's address space" is extended, for object introduction and deletion events, to include address spaces being employed on behalf of a remote user (or host). However, the focus remains on users in contrast to internal subjects as discussed in the DAC criterion. In addition, audit information must be stored in machine-readable form.

+ Rationale

For remote users, the network identifiers (e.g., internet address) can be used as identifiers of groups of individual users (e.g., "all users at Host A") to eliminate the maintenance that would be required if individual identification of remote users was employed. In this class (C2), however, it must be possible to identify (immediately or at some later time) the individuals represented by a group identifier. In all other respects, the interpretation is a straightforward extension of the criterion into the context of a network system.

3.1.3 Assurance

— — — — —

3.1.3.1 Operational Assurance

3.1.3.1.1 System Architecture

+ Statement from DoD 5200.28-STD

The TCB shall maintain a domain for its own execution that protects it from external interference or tampering (e.g., by modification of its code or data structures). Resources controlled by the TCB may be a defined subset of the subjects and objects in the ADP system. THE TCB SHALL MAINTAIN PROCESS ISOLATION THROUGH THE PROVISION OF DISTINCT ADDRESS SPACES UNDER ITS CONTROL. The TCB shall isolate the resources to be protected so that they are subject to the access control and auditing requirements.

+ Interpretation

The system architecture criterion must be met individually by all NTCB partitions. Implementation of the requirement that the NTCB maintain a domain for its own execution is achieved by having each NTCB partition maintain a domain for its own execution. SINCE EACH COMPONENT IS ITSELF A DISTINCT DOMAIN IN THE OVERALL NETWORK SYSTEM, THIS ALSO SATISFIES THE REQUIREMENT FOR PROCESS ISOLATION THROUGH DISTINCT ADDRESS SPACES IN THE SPECIAL CASE WHERE A COMPONENT HAS ONLY A SINGLE SUBJECT.

The subset of network resources over which the NTCB has control are the union of the sets of resources over which the NTCB partitions have control. Code and data structures belonging to the NTCB, transferred among NTCB subjects (i.e., subjects outside the reference monitor but inside the NTCB) belonging to different NTCB partitions, must be protected against external interference or tampering. For example, a cryptographic checksum or physical means may be employed to protect user authentication data exchanged between NTCB partitions.

Each NTCB partition provides isolation of RESOURCES (WITHIN ITS COMPONENT) TO BE PROTECTED IN accord with the network system architecture and security policy SO THAT "SUPPORTING ELEMENTS" (E.G., DAC AND USER IDENTIFICATION) FOR THE SECURITY MECHANISMS OF THE NETWORK SYSTEM ARE STRENGTHENED COMPARED TO C2, FROM AN ASSURANCE POINT OF VIEW, THROUGH THE PROVISION OF DISTINCT ADDRESS SPACES UNDER CONTROL OF THE NTCB.

AS DISCUSSED IN THE DISCRETIONARY ACCESS CONTROL SECTION, THE DAC MECHANISM OF A NTCB PARTITION MAY BE IMPLEMENTED AT THE INTERFACE OF THE REFERENCE MONITOR OR MAY BE DISTRIBUTED IN SUBJECTS THAT ARE PART OF THE NTCB IN THE SAME OR DIFFERENT COMPONENT. WHEN DISTRIBUTED IN NTCB SUBJECTS (I.E., WHEN OUTSIDE THE REFERENCE MONITOR), THE ASSURANCE REQUIREMENTS FOR THE DESIGN AND IMPLEMENTATION OF THE DAC SHALL BE THOSE OF CLASS C2 FOR ALL NETWORKS OF CLASS C2 OR ABOVE.

+ Rationale

The requirement for the protection of communications between NTCB partitions is specifically directed to subjects that are part of the NTCB partitions. Any requirements for such protection for the subjects that are outside the NTCB partitions are addressed in response to the integrity requirements of the security policy.

THE PROVISION OF DISTINCT ADDRESS SPACES UNDER THE CONTROL OF THE NTCB PROVIDES THE ABILITY TO SEPARATE SUBJECTS ACCORDING TO SENSITIVITY LEVEL. THIS REQUIREMENT IS INTRODUCED AT B1 SINCE IT IS AN ABSOLUTE NECESSITY IN ORDER TO IMPLEMENT MANDATORY ACCESS CONTROLS.

3.1.3.1.2 System Integrity

+ Statement from DoD 5200.28-STD

Hardware and/or software features shall be provided that can be used to periodically validate the correct operation of the on-site hardware and firmware elements of the TCB.

+ Interpretation

Implementation of the requirement is partly achieved by having hardware and/or software features that can be used to periodically validate the correct operation of the hardware and firmware elements of each component's NTCB partition. Features shall also be provided to validate the identity and correct operation of a component prior to its incorporation in the network system and throughout system operation. For example, a protocol could be designed that enables the components of the partitioned NTCB to exchange messages periodically and validate each other's correct response. The protocol shall be able to determine the remote entity's ability to respond. NTCB partitions shall provide the capability to report to network administrative personnel the failures detected in other NTCB partitions.

Intercomponent protocols implemented within a NTCB shall be designed in such a way as to provide correct operation in the case of failures of network communications or individual components. The allocation of MANDATORY AND discretionary access control policy in a network may require communication between trusted subjects that are part of the NTCB partitions in different components. This communication is normally implemented with a protocol between the subjects as peer entities. Incorrect access within a component shall not result from failure of an NTCB partition to communicate with other components.

+ Rationale

The first paragraph of the interpretation is a straightforward extension of the requirement into the context of a network system and partitioned NTCB as defined for these network criteria.

NTCB protocols should be robust enough so that they permit the system to operate correctly in the case of localized failure. The purpose of this protection is to preserve the integrity of the NTCB itself. It is not unusual for one or more components in a network to be inoperative at any time, so it is important to minimize the effects of such failures on the rest of the network. Additional integrity and denial of service issues are addressed in Part II.

3.1.3.2 Life-Cycle Assurance

3.1.3.2.1 Security Testing

+ Statement from DoD 5200.28-STD

The security mechanisms of the ADP system shall be tested and found to work as claimed in the system documentation. A TEAM OF INDIVIDUALS WHO THOROUGHLY UNDERSTAND THE SPECIFIC IMPLEMENTATION OF THE TCB SHALL SUBJECT ITS DESIGN DOCUMENTATION, SOURCE CODE, AND OBJECT CODE TO THROUGH ANALYSIS AND TESTING. THEIR OBJECTIVES SHALL BE: TO UNCOVER ALL DESIGN AND IMPLEMENTATION FLAWS THAT WOULD PERMIT A SUBJECT EXTERNAL TO THE TCB TO READ, CHANGE, OR DELETE DATA NORMALLY DENIED UNDER THE MANDATORY OR DISCRETIONARY SECURITY POLICY ENFORCED BY THE TCB; AS WELL AS TO ASSURE THAT NO SUBJECT (WITHOUT AUTHORIZATION TO DO SO) IS ABLE TO CAUSE THE TCB TO ENTER A STATE SUCH THAT IT IS UNABLE TO RESPOND TO COMMUNICATIONS INITIATED BY OTHER USERS. ALL DISCOVERED FLAWS SHALL BE REMOVED OR NEUTRALIZED AND THE TCB RETESTED TO DEMONSTRATE THAT THEY HAVE BEEN ELIMINATED AND THAT NEW FLAWS HAVE NOT BEEN INTRODUCED. (See the Security Testing Guidelines.)

+ Interpretation

Testing of a component will require a testbed that exercises the interfaces and protocols of the component including tests under exceptional conditions. The testing of a security mechanism of the network system for meeting this criterion shall be an integrated testing procedure involving all components containing an NTCB partition that implement the given mechanism. This integrated testing is additional to any individual component tests involved in the evaluation of the network system. The sponsor should iden-

tify the allowable set of configurations including the sizes of the networks. Analysis or testing procedures and tools shall be available to test the limits of these configurations. A change in configuration within the allowable set of configurations does not require retesting.

THE TESTING OF EACH COMPONENT WILL INCLUDE THE INTRODUCTION OF SUBJECTS EXTERNAL TO THE NTCB PARTITION FOR THE COMPONENT THAT WILL ATTEMPT TO READ, CHANGE, OR DELETE DATA NORMALLY DENIED. IF THE NORMAL INTERFACE TO THE COMPONENT DOES NOT PROVIDE A MEANS TO CREATE THE SUBJECTS NEEDED TO CONDUCT SUCH A TEST, THEN THIS PORTION OF THE TESTING SHALL USE A SPECIAL VERSION OF THE UNTRUSTED SOFTWARE FOR THE COMPONENT THAT RESULTS IN SUBJECTS THAT MAKE SUCH ATTEMPTS. THE RESULTS SHALL BE SAVED FOR TEST ANALYSIS. SUCH SPECIAL VERSIONS SHALL HAVE AN NTCB PARTITION THAT IS IDENTICAL TO THAT FOR THE NORMAL CONFIGURATION OF THE COMPONENT UNDER EVALUATION.

THE TESTING OF THE MANDATORY CONTROLS SHALL INCLUDE TESTS TO DEMONSTRATE THAT THE LABELS FOR INFORMATION IMPORTED AND/OR EXPORTED TO/FROM THE COMPONENT ACCURATELY REPRESENT THE LABELS MAINTAINED BY THE NTCB PARTITION FOR THE COMPONENT FOR USE AS THE BASIS FOR ITS MANDATORY ACCESS CONTROL DECISIONS. THE TESTS SHALL INCLUDE EACH TYPE OF DEVICE, WHETHER SINGLE-LEVEL OR MULTILEVEL, SUPPORTED BY THE COMPONENT.

+ Rationale

THE PHRASE "NO SUBJECT (WITHOUT AUTHORIZATION TO DO SO) IS ABLE TO CAUSE THE TCB TO ENTER A STATE SUCH THAT IT IS UNABLE TO RESPOND TO COMMUNICATIONS INITIATED BY OTHER USERS" RELATES TO THE SECURITY SERVICES (PART II OF THIS TNI) FOR THE DENIAL OF SERVICE PROBLEM, AND TO CORRECTNESS OF THE PROTOCOL IMPLEMENTATIONS.

Testing is AN IMPORTANT method available in this evaluation division to gain any assurance that the security mechanisms perform their intended function. A MAJOR PURPOSE OF TESTING IS TO DEMONSTRATE THE SYSTEM'S RESPONSE TO INPUTS TO THE NTCB PARTITION FROM UNTRUSTED (AND POSSIBLY MALICIOUS) SUBJECTS.

IN CONTRAST TO GENERAL PURPOSE SYSTEMS THAT ALLOW FOR THE DYNAMIC CREATION OF NEW PROGRAMS AND THE INTRODUCTIONS OF NEW PROCESSES (AND HENCE NEW SUBJECTS) WITH USER SPECIFIED SECURITY PROPERTIES, MANY NETWORK COMPONENTS HAVE NO METHOD FOR INTRODUCING NEW PROGRAMS AND/OR PROCESSES DURING THEIR NORMAL OPERATION. THEREFORE, THE PROGRAMS NECESSARY FOR THE TESTING MUST BE INTRODUCED AS SPECIAL VERSIONS OF THE SOFTWARE RATHER THAN AS THE RESULT OF NORMAL INPUTS BY THE TEST TEAM. HOWEVER, IT MUST BE INSURED THAT THE NTCB PARTITION USED FOR SUCH TESTS IS IDENTICAL TO THE ONE UNDER

EVALUATION.

SENSITIVITY LABELS SERVE A CRITICAL ROLE IN MAINTAINING THE SECURITY OF THE MANDATORY ACCESS CONTROLS IN THE NETWORK. ESPECIALLY IMPORTANT TO NETWORK SECURITY IS THE ROLE OF THE LABELS FOR INFORMATION COMMUNICATED BETWEEN COMPONENTS - EXPLICIT LABELS FOR MULTILEVEL DEVICES AND IMPLICIT LABELS FOR SINGLE-LEVEL DEVICES. THEREFORE THE TESTING FOR CORRECT LABELS IS HIGHLIGHTED.

3.1.3.2.2 Design Specification and Verification

+ Statement from DoD 5200.28-STD

AN INFORMAL OR FORMAL MODEL OF THE SECURITY POLICY SUPPORTED BY THE TCB SHALL BE MAINTAINED OVER THE LIFE CYCLE OF THE ADP SYSTEM AND DEMONSTRATED TO BE CONSISTENT WITH ITS AXIOMS.

+ Interpretation

THE OVERALL NETWORK SECURITY POLICY EXPRESSED IN THIS MODEL WILL PROVIDE THE BASIS FOR THE MANDATORY ACCESS CONTROL POLICY EXERCISED BY THE NTCB OVER SUBJECTS AND STORAGE OBJECTS IN THE ENTIRE NETWORK. THE POLICY WILL ALSO BE THE BASIS FOR THE DISCRETIONARY ACCESS CONTROL POLICY EXERCISED BY THE NTCB TO CONTROL ACCESS OF NAMED USERS TO NAMED OBJECTS. DATA INTEGRITY REQUIREMENTS ADDRESSING THE EFFECTS OF UNAUTHORIZED MSM NEED NOT BE INCLUDED IN THIS MODEL. THE OVERALL NETWORK POLICY MUST BE DECOMPOSED INTO POLICY ELEMENTS THAT ARE ALLOCATED TO APPROPRIATE COMPONENTS AND USED AS THE BASIS FOR THE SECURITY POLICY MODEL FOR THOSE COMPONENTS.

THE LEVEL OF ABSTRACTION OF THE MODEL, AND THE SET OF SUBJECTS AND OBJECTS THAT ARE EXPLICITLY REPRESENTED IN THE MODEL, WILL BE AFFECTED BY THE NTCB PARTITIONING. SUBJECTS AND OBJECTS MUST BE REPRESENTED EXPLICITLY IN THE MODEL FOR THE PARTITION IF THERE IS SOME NETWORK COMPONENT WHOSE NTCB PARTITION EXERCISES ACCESS CONTROL OVER THEM. THE MODEL SHALL BE STRUCTURED SO THAT THE AXIOMS AND ENTITIES APPLICABLE TO INDIVIDUAL NETWORK COMPONENTS ARE MANIFEST. GLOBAL NETWORK POLICY ELEMENTS THAT ARE ALLOCATED TO COMPONENTS SHALL BE REPRESENTED BY THE MODEL FOR THAT COMPONENT.

+ Rationale

THE TREATMENT OF THE MODEL DEPENDS TO A GREAT EXTENT ON THE DEGREE OF INTEGRATION OF THE COMMUNICATIONS SERVICE INTO A DISTRIBUTED SYSTEM. IN A CLOSELY COUPLED DISTRIBUTED SYSTEM, ONE MIGHT USE A MODEL THAT CLOSELY RESEMBLES ONE APPROPRIATE FOR A STAND-ALONE COMPUTER SYSTEM.

IN OTHER CASES, THE MODEL OF EACH PARTITION WILL BE EXPECTED TO SHOW THE ROLE OF THE NTCB PARTITION IN EACH KIND OF COMPONENT. IT WILL MOST LIKELY CLARIFY THE MODEL, ALTHOUGH NOT PART OF THE MODEL, TO SHOW ACCESS RESTRICTIONS IMPLIED BY THE SYSTEM DESIGN; FOR EXAMPLE, SUBJECTS REPRESENTING PROTOCOL ENTITIES MIGHT HAVE ACCESS ONLY TO OBJECTS CONTAINING DATA UNITS AT THE SAME LAYER OF PROTOCOL. THE ALLOCATION OF SUBJECTS AND OBJECTS TO DIFFERENT PROTOCOL LAYERS IS A PROTOCOL DESIGN CHOICE WHICH NEED NOT BE REFLECTED IN THE SECURITY POLICY MODEL.

3.1.4 Documentation.

3.1.4.1 Security Features User's Guide

+ Statement from DoD 5200.28-STD

A single summary, chapter, or manual in user documentation shall describe the protection mechanisms provided by the TCB, interpretations on their use, and how they interact with one another.

+ Interpretation

This user documentation describes user visible protection mechanisms at the global (network system) level and at the user interface of each component, and the interaction among these.

+ Rationale

The interpretation is an extension of the requirement into the context of a network system as defined for these network criteria. Documentation of protection mechanisms provided by individual components is required by the criteria for trusted computer systems that are applied as appropriate for the individual components.

3.1.4.2 Trusted Facility Manual

+ Statement from DoD 5200.28-STD

A manual addressed to the ADP system administrator shall present cautions about functions and privileges that should be controlled when running a secure facility. The procedures for examining and maintaining the audit files as well as the detailed audit record structure for each type of audit event shall be given. THE MANUAL SHALL DESCRIBE THE OPERATOR AND ADMINISTRATOR FUNCTIONS RELATED TO SECURITY, TO INCLUDE

CHANGING THE SECURITY CHARACTERISTICS OF A USER. IT SHALL PROVIDE INTERPRETATIONS ON THE CONSISTENT AND EFFECTIVE USE OF THE PROTECTION FEATURES OF THE SYSTEM, HOW THEY INTERACT, HOW TO SECURELY GENERATE A NEW TCB, AND FACILITY PROCEDURES, WARNINGS, AND PRIVILEGES THAT NEED TO BE CONTROLLED IN ORDER TO OPERATE THE FACILITY IN A SECURE MANNER.

+ Interpretation

This manual shall contain specifications and procedures to assist the system administrator(s) maintain cognizance of the network configuration. These specifications and procedures shall address the following:

1. The hardware configuration of the network itself;
2. The implications of attaching new components to the network;
3. The case where certain components may periodically leave the network (e.g., by crashing, or by being disconnected) and then rejoin;
4. Network configuration aspects that can impact the security of the network system; (For example, the manual should describe for the network system administrator the interconnections among components that are consistent with the overall network system architecture.)
5. Loading or modifying NTCB software or firmware (e.g., down-line loading).

The physical and administrative environmental controls shall be specified. Any assumptions about security of a given network should be clearly stated (e.g., the fact that all communications links must be physically protected to a certain level).

+ Rationale

There may be multiple system administrators with diverse responsibilities. The technical security measures described by these criteria must be used in conjunction with other forms of security in order to achieve security of the network. Additional forms include administrative security, physical security, emanations security, etc.

Extension of this criterion to cover configuration aspects of the network is needed because, for example, proper interconnection of components is typically essential to achieve a correct realization of the network architecture.

AS MENTIONED IN THE SECTION ON LABEL INTEGRITY, cryptography is one common mechanism employed to protect communication circuits. Encryption transforms the representation of information so that it is unintelligible to unauthorized subjects. Reflecting this transformation, the sensitivity of the ciphertext is generally lower than the cleartext. If encryption methodologies are employed, they shall be approved by the National Security Agency (NSA).

The encryption algorithm and its implementation are outside the scope of these interpretations. This algorithm and implementation may be implemented in a separate device or may be a function of a subject in a component not dedicated to encryption. Without prejudice, either implementation packaging is referred to as an encryption mechanism herein.

3.1.4.3 Test Documentation

+ Statement from DoD 5200.28-STD

The system developer shall provide to the evaluators a document that describes the test plan, test procedures that show how the security mechanisms were tested, and results of the security mechanisms' functional testing.

+ Interpretation

The "system developer" is interpreted as "the network system sponsor". The description of the test plan should establish the context in which the testing was or should be conducted. The description should identify any additional test components that are not part of the system being evaluated. This includes a description of the test-relevant functions of such test components and a description of the interfacing of those test components to the system being evaluated. The description of the test plan should also demonstrate that the tests adequately cover the network security policy. The tests should include the features described in the System Architecture and the System Integrity sections. The tests should also include network configuration and sizing.

+ Rationale

The entity being evaluated may be a networking subsystem (see Appendix A) to which other components must be added to make a complete network system. In that case, this interpretation is extended to include contextual definition because, at evaluation time, it is not possible to validate the test plans without the description of the context for testing the networking subsystem.

3.1.4.4 Design Documentation

+ Statement from DoD 5200.28-STD

Documentation shall be available that provides a description of the manufacturer's philosophy of protection and an explanation of how this philosophy is translated into the TCB. If the TCB is composed of distinct modules, the interfaces between these modules shall be described. AN INFORMAL OR FORMAL DESCRIPTION OF THE SECURITY POLICY MODEL ENFORCED BY THE TCB SHALL BE AVAILABLE AND AN EXPLANATION PROVIDED TO SHOW THAT IT IS SUFFICIENT TO ENFORCE THE SECURITY POLICY. THE SPECIFIC TCB PROTECTION MECHANISMS SHALL BE IDENTIFIED AND AN EXPLANATION GIVEN TO SHOW THAT THEY SATISFY THE MODEL.

+ Interpretation

Explanation of how the sponsor's philosophy of protection is translated into the NTCB shall include a description of how the NTCB is partitioned. The security policy also shall be stated. The description of the interfaces between the NTCB modules shall include the interface(s) between NTCB partitions and modules within the partitions if the modules exist. The sponsor shall describe the security architecture and design, including the allocation of security requirements among components. Appendix A addresses component evaluation issues.

AS STATED IN THE INTRODUCTION TO DIVISION B, THE SPONSOR MUST DEMONSTRATE THAT THE NTCB EMPLOYS THE REFERENCE MONITOR CONCEPT. THE SECURITY POLICY MODEL MUST BE A MODEL FOR A REFERENCE MONITOR.

THE SECURITY POLICY MODEL FOR EACH PARTITION IMPLEMENTING A REFERENCE MONITOR SHALL FULLY REPRESENT THE ACCESS CONTROL POLICY SUPPORTED BY THE PARTITION, INCLUDING THE DISCRETIONARY AND MANDATORY SECURITY POLICY FOR SECRECY AND/OR INTEGRITY. FOR THE MANDATORY POLICY THE SINGLE DOMINANCE RELATION FOR SENSITIVITY LABELS, INCLUDING SECRECY AND/OR INTEGRITY COMPONENTS, SHALL BE PRECISELY DEFINED.

+ Rationale

The interpretation is a straightforward extension of the requirement into the context of a network system as defined for this network interpretation. Other documentation, such as description of components and description of operating environment(s) in which the networking subsystem or network system is designed to function, is required elsewhere, e.g., in the Trusted Facility Manual.

In order to be evaluated, a network must possess a coherent Network Security Architecture and Design. (Interconnection of components that do not adhere to such a single coherent Network Security Architecture is addressed in the Interconnection of Accredited AIS, Appendix C.) The Network Security Architecture must address the security-relevant policies, objectives, and protocols. The Network Security Design specifies the interfaces and services that must be incorporated into the network so that it can be evaluated as a trusted entity. There may be multiple designs that conform to the same architecture but are more or less incompatible and non-interoperable (except through the Interconnection Rules). Security related mechanisms requiring cooperation among components are specified in the design in terms of their visible interfaces; mechanisms having no visible interfaces are not specified in this document but are left as implementation decisions.

The Network Security Architecture and Design must be available from the network sponsor before evaluation of the network, or any component, can be undertaken. The Network Security Architecture and Design must be sufficiently complete, unambiguous, and free from obvious flaws to permit the construction or assembly of a trusted network based on the structure it specifies.

When a component is being designed or presented for evaluation, or when a network assembled from components is assembled or presented for evaluation, there must be a priori evidence that the Network security Architecture and Design are satisfied. That is, the components can be assembled into a network that conforms in every way with the Network Security Architecture and Design to produce a physical realization that is trusted to the extent that its evaluation indicates.

In order for a trusted network to be constructed from components that can be built independently, the Network Security Architecture and Design must completely and unambiguously define the security functionality of components as well as the interfaces between or among components. The Network Security Architecture and Design must be evaluated to determine that a network constructed to its specifications will in fact be trusted, that is, it will be evaluable under these interpretations.

THE TERM "MODEL" IS USED IN SEVERAL DIFFERENT WAYS IN A NETWORK CONTEXT, E.G., A "PROTOCOL REFERENCE MODEL," A "FORMAL NETWORK MODEL," ETC. ONLY THE "SECURITY POLICY MODEL" IS ADDRESSED BY THIS REQUIREMENT AND IS SPECIFICALLY INTENDED TO MODEL THE INTERFACE, VIZ., "SECURITY PERIMETER," OF THE REFERENCE MONITOR AND MUST MEET ALL THE REQUIREMENTS DEFINED IN THE TCSEC. IT MUST BE SHOWN THAT ALL PARTS OF THE TCB

ARE A VALID INTERPRETATION OF THE SECURITY POLICY MODEL, I.E., THAT THERE IS NO CHANGE TO THE SECURE STATE EXCEPT AS REPRESENTED BY THE MODEL.

3.2 CLASS (B2): STRUCTURED PROTECTION

IN CLASS (B2) NETWORK SYSTEMS, THE NTCB IS BASED ON A CLEARLY DEFINED AND DOCUMENTED FORMAL SECURITY POLICY MODEL THAT REQUIRES THE DISCRETIONARY AND MANDATORY ACCESS CONTROL ENFORCEMENT FOUND IN CLASS (B1) NETWORK SYSTEMS TO BE EXTENDED TO ALL SUBJECTS AND OBJECTS IN THE NETWORK SYSTEM. IN ADDITION, COVERT CHANNELS ARE ADDRESSED. THE NTCB MUST BE CAREFULLY STRUCTURED INTO PROTECTION-CRITICAL AND NON-PROTECTION-CRITICAL ELEMENTS. THE NTCB INTERFACE IS WELL-DEFINED, AND THE NTCB DESIGN AND IMPLEMENTATION ENABLE IT TO BE SUBJECTED TO MORE THOROUGH TESTING AND MORE COMPLETE REVIEW. AUTHENTICATION MECHANISMS ARE STRENGTHENED, TRUSTED FACILITY MANAGEMENT IS PROVIDED IN THE FORM OF SUPPORT FOR SYSTEM ADMINISTRATOR AND OPERATOR FUNCTIONS, AND STRINGENT CONFIGURATION MANAGEMENT CONTROLS ARE IMPOSED. THE SYSTEM IS RELATIVELY RESISTANT TO PENETRATION. THE FOLLOWING ARE MINIMAL REQUIREMENTS FOR SYSTEM ASSIGNED A CLASS (B2) RATING.

3.2.1 Security Policy

+ Statement from DoD 5200.28-STD

Implied from the Introduction to the TCSEC.

+ Interpretation

The network sponsor shall describe the overall network security policy enforced by the NTCB. At a minimum, this policy shall include the discretionary and mandatory requirements applicable to this class. The policy may require data secrecy, or data integrity, or both. The policy is an access control policy having two primary components: mandatory and discretionary. The policy shall include a discretionary policy for protecting the information being processed based on the authorizations of individuals, users, or groups of users. This access control policy statement shall describe the requirements on the network to prevent or detect "reading or destroying" sensitive information by unauthorized users or errors. The mandatory policy must define the set of distinct sensitivity levels

that it supports. For the Class B1 or above the mandatory policy shall be based on the labels associated with the information that reflects its sensitivity with respect to secrecy and/or integrity, where applicable, and labels associated with users to reflect their authorization to access such information. Unauthorized users include both those that are not authorized to use the network at all (e.g., a user attempting to use a passive or active wire tap) or a legitimate user of the network who is not authorized to access a specific piece of information being protected.

Note that "users" does not include "operators," "system programmers," "technical control officers," "system security officers," and other system support personnel. They are distinct from users and are subject to the Trusted Facility Manual and the System Architecture requirements. Such individuals may change the system parameters of the network system, for example, by defining membership of a group. These individuals may also have the separate role of users.

SECURITY POLICY: The network sponsor shall define the form of the discretionary and mandatory security policy that is enforced in the network to prevent unauthorized users from reading the sensitive information entrusted to the network.

DATA INTEGRITY POLICY: The network sponsor shall define the discretionary and mandatory integrity policy to prevent unauthorized users from modifying, viz., writing, sensitive information. The definition of data integrity presented by the network sponsor refers to the requirement that the information has not been subjected to unauthorized modification in the network. The mandatory integrity policy enforced by the NTCB cannot, in general, prevent modification while information is being transmitted between components. However, an integrity sensitivity label may reflect the confidence that the information has not been subjected to transmission errors because of the protection afforded during transmission. This requirement is distinct from the requirement for label integrity.

+ Rationale

The word "sponsor" is used in place of alternatives (such as "vendor," "architect," "manufacturer," and "developer") because the alternatives indicate people who may not be available, involved, or relevant at the time that a network system is proposed for evaluation.

A trusted network is able to control both the reading

and writing of shared sensitive information. Control of writing is used to protect against destruction of information. A network normally is expected to have policy requirements to protect both the secrecy and integrity of the information entrusted to it. In a network the integrity is frequently as important or more important than the secrecy requirements. Therefore the secrecy and/or integrity policy to be enforced by the network must be stated for each network regardless of its evaluation class. The assurance that the policy is faithfully enforced is reflected in the evaluation class of the network.

This control over modification is typically used to protect information so that it may be relied upon and to control the potential harm that would result if the information were corrupted. The overall network policy requirements for integrity includes the protection for data both while being processed in a component and while being transmitted in the network. The access control policy enforced by the NTCB relates to the access of subjects to objects within each component. Communications integrity addressed within Part II relates to information while being transmitted.

The mandatory integrity policy (at class B1 and above) in some architectures may be useful in supporting the linkage between the connection oriented abstraction introduced in the Introduction and the individual components of the network. For example, in a key distribution center for end-to-end encryption, a distinct integrity category may be assigned to isolate the key generation code and data from possible modification by other supporting processes in the same component, such as operator interfaces and audit.

The mandatory integrity policy for some architecture may define an integrity sensitivity label that reflects the specific requirements for ensuring that information has not been subject to random errors in excess of a stated limit nor to unauthorized message stream modification (MSM) -. The specific metric associated with an integrity sensitivity label will generally reflect the intended applications of the network.

3.2.1.1 Discretionary Access Control

+ Statement from DoD 5200.28-STD

The TCB shall define and control access between named users and named objects (e.g., files and programs) in the ADP system. The enforcement mechanism (e.g., self/group/public

- See Voydock, Victor L. and Stephen T. Kent, "Security Mechanisms in High-Level Network Protocols," Com-

controls, access control lists) shall allow users to specify and control sharing of those objects by named individuals or defined groups of individuals, or both, and shall provide controls to limit propagation of access rights. The discretionary access control mechanism shall, either by explicit user action or by default, provide that objects are protected from unauthorized access. These access controls shall be capable of including or excluding access to the granularity of a single user. Access permission to an object by users not already possessing access permission shall only be assigned by authorized users.

+ Interpretation

The discretionary access control (DAC) mechanism(s) may be distributed over the partitioned NTCB in various ways. Some part, all, or none of the DAC may be implemented in a given component of the network system. In particular, components that support only internal subjects (i.e., that have no subjects acting as direct surrogates for users), such as a public network packet switch, might not implement the DAC mechanism(s) directly (e.g., they are unlikely to contain access control lists).

Identification of users by groups may be achieved in various ways in the networking environment. For example, the network identifiers (e.g., internet addresses) for various components (e.g., hosts, gateways) can be used as identifiers of groups of individual users (e.g., "all users at Host A," "all users of network Q") so long as the individuals involved in the group are implied by the group identifier. For example, Host A might employ a particular group-id, for which it maintains a list of explicit users in that group, in its network exchange with Host B, which accepts the group-id under the conditions of this interpretation.

For networks, individual hosts will impose need-to-know controls over their users on the basis of named individuals - much like (in fact, probably the same) controls used when there is no network connection.

When group identifiers are acceptable for access control, the identifier of some other host may be employed, to eliminate the maintenance that would be required if individual identification of remote users was employed. In class C2 and higher, however, it must be possible from that audit record to identify (immediately or at some later time) exactly the individuals represented by a group identifier at the time of the use of that identifier. There is allowed to

be an uncertainty because of elapsed time between changes in the group membership and the enforcement in the access control mechanisms.

The DAC mechanism of a NTCB partition may be implemented at the interface of the reference monitor or may be distributed in subjects that are part of the NTCB in the same or different component. The reference monitor manages all the physical resources of the system and from them creates the abstraction of subjects and objects that it controls. Some of these subjects and objects may be used to implement a part of the NTCB. When the DAC mechanism is distributed in such NTCB subjects (i.e., when outside the reference monitor), the assurance requirements (see the Assurance section) for the design and implementation of the DAC shall be those of class C2 for all networks of class C2 or above.

When integrity is included as part of the network discretionary security policy, the above interpretations shall be specifically applied to the controls over modification, viz, the write mode of access, within each component based on identified users or groups of users.

+ Rationale

In this class, the supporting elements of the overall DAC mechanism are required to isolate information (objects) that supports DAC so that it is subject to auditing requirements (see the System Architecture section). The use of network identifiers to identify groups of individual users could be implemented, for example, as an X.25 community of interest in the network protocol layer (layer 3). In all other respects, the supporting elements of the overall DAC mechanism are treated exactly as untrusted subjects are treated with respect to DAC in an ADP system, with the same result as noted in the interpretation.

A typical situation for DAC is that a surrogate process for a remote user will be created in some host for access to objects under the control of the NTCB partition within that host. The interpretation requires that a user identifier be assigned and maintained for each such process by the NTCB, so that access by a surrogate process is subject to essentially the same discretionary controls as access by a process acting on behalf of a local user would be. However, within this interpretation a range of possible interpretations of the assigned user identification is permitted.

The most obvious situation would exist if a global database of network users were to be made permanently available on demand to every host, (i.e., a name server existed) so that all user identifications were globally meaningful.

It is also acceptable, however, for some NTCB partitions to maintain a database of locally-registered users for its own use. In such a case, one could choose to inhibit the creation of surrogate processes for locally unregistered users, or (if permitted by the local policy) alternatively, to permit the creation of surrogate processes with preselected user and group identifiers which, in effect, identify the process as executing on behalf of a member of a group of users on a particular remote host. The intent of the words concerning audit in the interpretation is to provide a minimally acceptable degree of auditability for cases such as the last described. What is required is that there be a capability, using the audit facilities provided by the network NTCB partitions involved, to determine who was logged in at the actual host of the group of remote users at the time the surrogate processing occurred.

Associating the proper user id with a surrogate process is the job of identification and authentication. This means that DAC is applied locally, with respect to the user id of the surrogate process. The transmission of the data back across the network to the user's host, and the creation of a copy of the data there, is not the business of DAC.

Components that support only internal subjects impact the implementation of the DAC by providing services by which information (e.g., a user-id) is made available to a component that makes a DAC decision. An example of the latter would be the case that a user at Host A attempts to access a file at Host B. The DAC decision might be (and usually would be) made at Host B on the basis of a user-id transmitted from Host A to Host B.

Unique user identification may be achieved by a variety of mechanisms, including (a) a requirement for unique identification and authentication on the host where access takes place; (b) recognition of fully qualified network addresses authenticated by another host and forwarded to the host where access takes place; or (c) administrative support of a network-wide unique personnel identifier that could be authenticated and forwarded by another host as in (b) above, or could be authenticated and forwarded by a dedicated network identification and authentication server. The protocols which implement (b) or (c) are subject to the System Architecture requirements.

Network support for DAC might be handled in other ways than that described as "typical" above. In particular, some form of centralized access control is often proposed. An access control center may make all decisions for DAC, or it may share the burden with the hosts by controlling host-to-host connections, and leaving the hosts to decide on access to their objects by users at a limited set of remote hosts. In this case the access control center provides the linkage

between the connection oriented abstraction (as discussed in the Introduction) and the overall network security policy for DAC. In all cases the enforcement of the decision must be provided by the host where the object resides.

There are two forms of distribution for the DAC mechanism: implementing portions of the DAC in separate components, and supporting the DAC in subjects contained within the NTCB partition in a component. Since "the ADP system" is understood to be "the computer network" as a whole, each network component is responsible for enforcing security in the mechanisms allocated to it to ensure secure implementation of the network security policy. For traditional host systems it is frequently easy to also enforce the DAC along with the MAC within the reference monitor, per se, although a few approaches, such as virtual machine monitors, support DAC outside this interface.

In contrast to the universally rigid structure of mandatory policies (see the Mandatory Access Control section), DAC policies tend to be very network and system specific, with features that reflect the natural use of the system. For networks it is common that individual hosts will impose controls over their local users on the basis of named individuals-much like the controls used when there is no network connection. However, it is difficult to manage in a centralized manner all the individuals using a large network. Therefore, users on other hosts are commonly grouped together so that the controls required by the network DAC policy are actually based on the identity of the hosts or other components. A gateway is an example of such a component.

The assurance requirements are at the very heart of the concept of a trusted system. It is the assurance that determines if a system or network is appropriate for a given environment, as reflected, for example, in the Environments Guideline-. In the case of monolithic systems that have DAC integral to the reference monitor, the assurance requirements for DAC are inseparable from those of the rest of the reference monitor. For networks there is typically a much clearer distinction due to distributed DAC. The rationale for making the distinction in this network interpretation is that if major trusted network components can be made significantly easier to design and implement without reducing the ability to meet security policy, then trusted networks will be more easily available.

3.2.1.2 Object Reuse

+ Statement from DoD 5200.28-STD

All authorizations to the information contained within a storage object shall be revoked prior to initial assignment,

allocation or reallocation to a subject from the TCB's pool of unused storage objects. No information, including encrypted representations of information, produced by a prior subject's actions is to be available to any subject that obtains access to an object that has been released back to the system.

- Guidance for Applying the Department of Defense
Trusted Computer System Evaluation Criteria in Specific
Environments, CSC-STD-003-85.

+ Interpretation

The NTCB shall ensure that any storage objects that it controls (e.g., message buffers under the control of a NTCB partition in a component) contain no information for which a subject in that component is not authorized before granting access. This requirement must be enforced by each of the NTCB partitions.

+ Rationale

In a network system, storage objects of interest are things that the NTCB directly controls, such as message buffers in components. Each component of the network system must enforce the object reuse requirement with respect to the storage objects of interest as determined by the network security policy. For example, the DAC requirement in this division leads to the requirement here that message buffers be under the control of the NTCB partition. A buffer assigned to an internal subject may be reused at the discretion of that subject which is responsible for preserving the integrity of message streams. Such controlled objects may be implemented in physical resources, such as buffers, disk sectors, tape space, and main memory, in components such as network switches.

3.2.1.3 Labels

+ Statement from DoD 5200.28-STD

Sensitivity labels associated with each ADP SYSTEM RESOURCE (E.G., SUBJECT, STORAGE OBJECT, ROM) THAT IS DIRECTLY OR INDIRECTLY ACCESSIBLE BY SUBJECTS EXTERNAL TO THE TCB shall be maintained by the TCB. These labels shall be used as the basis for mandatory access control decisions. In order to import non-labeled data, the TCB shall request and receive

from an authorized user the sensitivity level of the data, and all such actions shall be auditable by the TCB.

+ Interpretation

Non-labeled data imported under the control of the NTCB partition will be assigned a label constrained by THE DEVICE LABELS OF the single-level device used to import it. Labels may include secrecy and integrity- components in accordance with the overall network security policy described by the network sponsor. Whenever the term "label" is used throughout this interpretation, it is understood to include both components as applicable. Similarly, the terms "single-level" and "multilevel" are understood to be based

- See, for example, Biba, K.J., "Integrity Consideration for Secure Computer Systems," ESD-TR-76-372, MTR-3153, The MITRE Corporation, Bedford, MA, April 1977.

on both the secrecy and integrity components of the policy. The mandatory integrity policy will typically have requirements, such as the probability of undetected message stream modification, that will be reflected in the label for the data so protected. For example, when data is imported its integrity label may be assigned based on mechanisms, such as cryptography, used to provide the assurance required by the policy. The NTCB shall assure that such mechanism are protected from tampering and are always invoked when they are the basis for a label.

IF THE SECURITY POLICY INCLUDES AN INTEGRITY POLICY, ALL ACTIVITIES THAT RESULT IN MESSAGE-STREAM MODIFICATION DURING TRANSMISSION ARE REGARDED AS UNAUTHORIZED ACCESSES IN VIOLATION OF THE INTEGRITY POLICY. THE NTCB SHALL HAVE AN AUTOMATED CAPABILITY FOR TESTING, DETECTING, AND REPORTING THOSE ERRORS/CORRUPTIONS THAT EXCEED SPECIFIED NETWORK INTEGRITY POLICY REQUIREMENTS. MESSAGE-STREAM MODIFICATION (MSM) COUNTERMEASURES SHALL BE IDENTIFIED. A TECHNOLOGY OF ADEQUATE STRENGTH SHALL BE SELECTED TO RESIST MSM. IF ENCRYPTION METHODOLOGIES ARE EMPLOYED, THEY SHALL BE APPROVED BY THE NATIONAL SECURITY AGENCY.

ALL OBJECTS MUST BE LABELED WITHIN EACH COMPONENT OF THE NETWORK THAT IS TRUSTED TO MAINTAIN SEPARATION OF MULTIPLE LEVELS OF INFORMATION. THE LABEL ASSOCIATED WITH ANY OBJECTS ASSOCIATED WITH SINGLE-LEVEL COMPONENTS WILL BE IDENTICAL TO THE LEVEL OF THAT COMPONENT. OBJECTS USED TO STORE NETWORK CONTROL INFORMATION, AND OTHER NETWORK STRUCTURES, SUCH AS ROUTING TABLES, MUST BE LABELED TO PREVENT UNAUTHORIZED ACCESS AND/OR MODIFICATION.

+ Rationale

The interpretation is an extension of the requirement into the context of a network system and partitioned NTCB as defined for these network interpretations. A single-level device may be regarded either as a subject or an object. A multilevel device is regarded as a trusted subject in which the security range of the subject is the minimum-maximum range of the data expected to be transmitted over the device.

The sensitivity labels for either secrecy or integrity or both may reflect non-hierarchical categories or hierarchical classification or both.

FOR A NETWORK IT IS NECESSARY THAT THIS REQUIREMENT BE APPLIED TO ALL NETWORK SYSTEM RESOURCES AT THE (B2) LEVEL AND ABOVE.

THE NTCB IS RESPONSIBLE FOR IMPLEMENTING THE NETWORK INTEGRITY POLICY, WHEN ONE EXISTS. THE NTCB MUST ENFORCE THAT POLICY BY ENSURING THAT INFORMATION IS ACCURATELY TRANSMITTED FROM SOURCE TO DESTINATION (REGARDLESS OF THE NUMBER OF INTERVENING CONNECTING POINTS). THE NTCB MUST BE ABLE TO COUNTER EQUIPMENT FAILURE, ENVIRONMENTAL DISRUPTIONS, AND ACTIONS BY PERSONS AND PROCESSES NOT AUTHORIZED TO ALTER THE DATA. PROTOCOLS THAT PERFORM CODE OR FORMAT CONVERSION SHALL PRESERVE THE INTEGRITY OF DATA AND CONTROL INFORMATION.

THE PROBABILITY OF AN UNDETECTED TRANSMISSION ERROR MAY BE SPECIFIED AS PART OF THE NETWORK SECURITY POLICY SO THAT THE ACCEPTABILITY OF THE NETWORK FOR ITS INTENDED APPLICATION MAY BE DETERMINED. THE SPECIFIC METRICS (E.G., PROBABILITY OF UNDETECTED MODIFICATION) SATISFIED BY THE DATA CAN BE REFLECTED IN THE INTEGRITY SENSITIVITY LABEL ASSOCIATED WITH THE DATA WHILE IT IS PROCESSED WITHIN A COMPONENT. IT IS RECOGNIZED THAT DIFFERENT APPLICATIONS AND OPERATIONAL ENVIRONMENTS (E.G., CRISIS AS COMPARED TO LOGISTIC) WILL HAVE DIFFERENT INTEGRITY REQUIREMENTS.

THE NETWORK SHALL ALSO HAVE AN AUTOMATED CAPABILITY OF TESTING FOR, DETECTING, AND REPORTING ERRORS THAT EXCEED A THRESHOLD CONSISTENT WITH THE OPERATIONAL MODE REQUIREMENTS. THE EFFECTIVENESS OF INTEGRITY COUNTERMEASURES MUST BE ESTABLISHED WITH THE SAME RIGOR AS THE OTHER SECURITY-RELEVANT PROPERTIES SUCH AS SECRECY.

CRYPTOGRAPHY IS OFTEN UTILIZED AS A BASIS TO PROVIDE DATA INTEGRITY ASSURANCE. MECHANISMS, SUCH AS MANIPULATION DETECTION CODES (MDC)-, MAY BE USED. THE ADEQUACY OF THE ENCRYPTION OR MDC ALGORITHM, THE CORRECTNESS OF THE PROTOCOL LOGIC, AND THE ADEQUACY OF IMPLEMENTATION MUST BE ESTABLISHED IN MSM COUNTERMEASURES DESIGN.

3.2.1.3.1 Label Integrity

+ Statement from DoD 5200.28-STD

Sensitivity labels shall accurately represent sensitivity levels of the specific subjects or objects with which they are associated. When exported by the TCB, sensitivity labels shall accurately and unambiguously represent the internal labels and shall be associated with the information being exported.

+ Interpretation

The phrase "exported by the TCB" is understood to include transmission of information from an object in one component to an object in another component. Information transferred between NTCB partitions is addressed in the System Integrity Section. The form of internal and external (exported) sensitivity labels may differ, but the meaning shall be the same. The NTCB shall, in addition, ensure that

- See Jueneman, R. R., "Electronic Document Authentication," IEEE Network Magazine, April 1987, pp 17-23.

correct association of sensitivity labels with the information being transported across the network is preserved.

As mentioned in the Trusted Facility Manual Section, encryption transforms the representation of information so that it is unintelligible to unauthorized subjects. Reflecting this transformation, the sensitivity level of the ciphertext is generally lower than the cleartext. It follows that cleartext and ciphertext are contained in different objects, each possessing its own label. The label of the cleartext must be preserved and associated with the ciphertext so that it can be restored when the cleartext is subsequently obtained by decrypting the ciphertext. If the cleartext is associated with a single-level device, the label of that cleartext may be implicit. The label may also be implicit in the key.

When information is exported to an environment where it is subject to deliberate or accidental modification, the TCB shall support the means, such as cryptographic checksums, to assure the accuracy of the labels. When there is a mandatory integrity policy, the policy will define the meaning of integrity labels.

+ Rationale

Encryption algorithms and their implementation are out-

side the scope of these interpretations. Such algorithms may be implemented in a separate device or may be incorporated in a subject of a larger component. Without prejudice, either implementation packaging is referred to as an encryption mechanism herein. If encryption methodologies are employed in this regard, they shall be approved by the National Security Agency (NSA). The encryption process is part of the Network Trusted Computer Base partition in the components in which it is implemented.

The encryption mechanism is not necessarily a multilevel device or multilevel subject, as these terms are used in these criteria. The process of encryption is multilevel by definition. The cleartext and ciphertext interfaces carry information of different sensitivity. An encryption mechanism does not process data in the sense of performing logical or arithmetic operations on that data with the intent of producing new data. The cleartext and ciphertext interfaces on the encryption mechanism must be separately identified as being single-level or multilevel. If the interface is single-level, then the sensitivity of the data is established by a trusted individual and implicitly associated with the interface; the Exportation to Single-Level Devices criterion applies.

If the interface is multilevel, then the data must be labeled; the Exportation to Multilevel Devices criterion applies. The network architect is free to select an acceptable mechanism for associating a label with an object. With reference to encrypted objects, the following examples are possible:

1. Include a label field in the protocol definition of the object.
2. Implicitly associate the label with the object through the encryption key. That is, the encryption key uniquely identifies a sensitivity level. A single or private key must be protected at the level of the data that it encrypts.

3.2.1.3.2 Exportation of Labeled Information

+ Statement from DoD 5200.28-STD

The TCB shall designate each communication channel and I/O device as either single-level or multilevel. Any change in this designation shall be done manually and shall be auditable by the TCB. The TCB shall maintain and be able to audit any change in the sensitivity level or levels associated with a communications channel or I/O device.

+ Interpretation

Each communication channel and network component shall be designated as either single-level or multilevel. Any change in this designation shall be done with the cognizance and approval of the administrator or security officer in charge of the affected components and the administrator or security officer in charge of the NTCB. This change shall be auditable by the network. The NTCB shall maintain and be able to audit any change in the DEVICE LABELS associated with a single-level communication channel or the range associated with a multilevel communication channel or component. The NTCB shall also be able to audit any change in the set of sensitivity levels associated with the information which can be transmitted over a multilevel communication channel or component.

+ Rationale

Communication channels and components in a network are analogous to communication channels and I/O devices in stand-alone systems. They must be designated as either multilevel (i.e., able to distinguish and maintain separation among information of various sensitivity levels) or single-level. As in the TCSEC, single-level devices may only be attached to single-level channels.

The level or set of levels of information that can be sent to a component or over a communication channel shall only change with the knowledge and approval of the security officers (or system administrator, if there is no security officer) of the network, and of the affected components. This requirement ensures that no significant security-relevant changes are made without the approval of all affected parties.

3.2.1.3.2.1 Exportation to Multilevel Devices

+ Statement from DoD 5200.28-STD

When the TCB exports an object to a multilevel I/O device, the sensitivity label associated with that object shall also be exported and shall reside on the same physical medium as the exported information and shall be in the same form (i.e., machine-readable or human-readable form). When the TCB exports or imports an object over a multilevel communications channel, the protocol used on that channel shall provide for the unambiguous pairing between the sensitivity labels and the associated information that is sent or received.

+ Interpretation

The components, including hosts, of a network shall be

interconnected over "multilevel communication channels," multiple single-level communication channels, or both, whenever the information is to be protected at more than a single sensitivity level. The protocol for associating the sensitivity label and the exported information shall provide the only information needed to correctly associate a sensitivity level with the exported information transferred over the multilevel channel between the NTCB partitions in individual components. This protocol definition must specify the representation and semantics of the sensitivity labels (i.e., the machine-readable label must uniquely represent the sensitivity level).

The "unambiguous" association of the sensitivity level with the communicated information shall meet the same level of accuracy as that required for any other label within the NTCB, as specified in the criterion for Label Integrity. This may be provided by protected and highly reliable direct physical layer connections, or by traditional cryptographic link protection in which any errors during transmission can be readily detected, or by use of a separate channel. THE RANGE OF INFORMATION IMPORTED OR EXPORTED MUST BE CONSTRAINED BY THE ASSOCIATED DEVICE LABELS.

+ Rationale

This protocol must specify the representation and semantics of the sensitivity labels. See the Mandatory Access Control Policies section in Appendix B. The multilevel device interface to (untrusted) subjects may be implemented either by the interface of the reference monitor, per se, or by a multilevel subject (e.g., a "trusted subject" as defined in the Bell-LaPadula Model) that provides the labels based on the internal labels of the NTCB partition.

The current state of the art limits the support for mandatory policy that is practical for secure networks. Reference monitor support to ensure the control over all the operations of each subject in the network must be completely provided within the single NTCB partition on which that subject interfaces to the NTCB. This means that the entire portion of the "secure state" represented in the FORMAL security policy model that may be changed by transitions invoked by this subject must be contained in the same component.

The secure state of an NTCB partition may be affected by events external to the component in which the NTCB partition resides (e.g., arrival of a message). The effect occurs asynchronously after being initiated by an event in another component or partition. For example, indeterminate delays may occur between the initiation of a message in one component, the arrival of the message in the NTCB partition

in another component, and the corresponding change to the secure state of the second component. Since each component is executing concurrently, to do otherwise would require some sort of network-wide control to synchronize state transitions, such as a global network-wide clock for all processors; in general, such designs are not practical and probably not even desirable. Therefore, the interaction between NTCB partitions is restricted to just communications between pairs (at least logically) of devices-multilevel devices if the device(s) can send/receive data of more than a single level. For broadcast channels the pairs are the sender and intended receiver(s). However, if the broadcast channel carries multiple levels of information, additional mechanism (e.g., cryptchecksum maintained by the TCB) may be required to enforce separation and proper delivery.

A common representation for sensitivity labels is needed in the protocol used on that channel and understood by both the sender and receiver when two multilevel devices (in this case, in two different components) are interconnected. Each distinct sensitivity level of the overall network policy must be represented uniquely in these labels.

Within a monolithic TCB, the accuracy of the sensitivity labels is generally assured by simple techniques, e.g., very reliable connections over very short physical connections, such as on a single printed circuit board or over an internal bus. In many network environments there is a much higher probability of accidentally or maliciously introduced errors, and these must be protected against.

3.2.1.3.2.2 Exportation to Single-Level Devices

+ Statement from DoD 5200.28-STD

Single-level I/O devices and single-level communication channels are not required to maintain the sensitivity labels of the information they process. However, the TCB shall include a mechanism by which the TCB and an authorized user reliably communicate to designate the single sensitivity level of information imported or exported via single-level communication channels or I/O devices.

+ Interpretation

Whenever one or both of two directly connected components is not trusted to maintain the separation of information of different sensitivity levels, or whenever the two directly connected components have only a single sensitivity level in common, the two components of the network shall communicate over a single-level channel. Single-level components and single-level communication channels are not required to maintain the sensitivity labels of the informa-

tion they process. However, the NTCB shall include a reliable communication mechanism by which the NTCB and an authorized user (VIA A TRUSTED PATH) or a subject within an NTCB partition can designate the single sensitivity level of information imported or exported via single-level communication channels or network components. THE LEVEL OF INFORMATION COMMUNICATED MUST EQUAL THE DEVICE LEVEL.

+ Rationale

Single-level communications channels and single-level components in networks are analogous to single level channels and I/O devices in stand-alone systems in that they are not trusted to maintain the separation of information of different sensitivity levels. The labels associated with data transmitted over those channels and by those components are therefore implicit; the NTCB associates labels with the data because of the channel or component, not because of an explicit part of the bit stream. Note that the sensitivity level of encrypted information is the level of the ciphertext rather than the original level(s) of the plaintext.

3.2.1.3.2.3 Labeling Human-Readable Output

+ Statement from DoD 5200.28-STD

The ADP system administrator shall be able to specify the printable label names associated with exported sensitivity labels. The TCB shall mark the beginning and end of all human-readable, paged, hardcopy output (e.g., line printer output) with human-readable sensitivity labels that properly¹ represent the sensitivity of the output. The TCB shall, by default, mark the top and bottom of each page of human-readable, paged, hardcopy output (e.g., line printer output) with human-readable sensitivity labels that properly¹ represent the sensitivity of the page. The TCB shall, by default and in an appropriate manner, mark other forms of human readable output (e.g., maps, graphics) with human-readable sensitivity labels that properly¹ represent the sensitivity of the output. Any override of these markings defaults shall be auditable by the TCB.

¹ The hierarchical classification component in human-readable sensitivity labels shall be equal to the greatest hierarchical classification of any of the information in the output that the labels refer to; the non-hierarchical category component shall include all of the non-hierarchical categories of the information in the output the labels refer to, but to no other non-hierarchical categories.

+ Interpretation

This criterion imposes no requirement to a component that produces no human-readable output. For those that do produce human-readable output, each sensitivity level that is defined to the network shall have a uniform meaning across all components. The network administrator, in conjunction with any affected component administrator, shall be able to specify the human-readable label that is associated with each defined sensitivity level.

+ Rationale

The interpretation is a straightforward extension of the requirement into the context of a network system and partitioned NTCB as defined for these network interpretations.

3.2.1.3.3 Subject Sensitivity Labels

+ Statement from DoD 5200.28-STD

THE TCB SHALL IMMEDIATELY NOTIFY A TERMINAL USER OF EACH CHANGE IN THE SENSITIVITY LEVEL ASSOCIATED WITH THAT USER DURING AN INTERACTIVE SESSION. A TERMINAL USER SHALL BE ABLE TO QUERY THE TCB AS DESIRED FOR A DISPLAY OF THE SUBJECT'S COMPLETE SENSITIVITY LABEL.

+ Interpretation

AN NTCB PARTITION SHALL IMMEDIATELY NOTIFY A TERMINAL USER ATTACHED TO ITS COMPONENT OF EACH CHANGE IN THE SENSITIVITY LEVEL ASSOCIATED WITH THAT USER.

+ Rationale

THE LOCAL NTCB PARTITION MUST ENSURE THAT THE USER UNDERSTANDS THE SENSITIVITY LEVEL OF INFORMATION SENT TO AND FROM A TERMINAL. WHEN A USER HAS A SURROGATE PROCESS IN ANOTHER COMPONENT, ADJUSTMENTS TO ITS LEVEL MAY OCCUR TO MAINTAIN COMMUNICATION WITH THE USER. THESE CHANGES MAY OCCUR ASYNCHRONOUSLY. SUCH ADJUSTMENTS ARE NECESSITATED BY MANDATORY ACCESS CONTROL AS APPLIED TO THE OBJECTS INVOLVED IN THE COMMUNICATION PATH.

3.2.1.3.4 Device Labels

+ Statement from DoD 5200.28-STD

THE TCB SHALL SUPPORT THE ASSIGNMENT OF MINIMUM AND MAXIMUM SENSITIVITY LEVELS TO ALL ATTACHED PHYSICAL DEVICES. THESE SENSITIVITY LEVELS SHALL BE USED BY THE TCB TO ENFORCE CONSTRAINTS IMPOSED BY THE PHYSICAL ENVIRONMENTS IN WHICH THE DEVICES ARE LOCATED.

+ Interpretation

THIS REQUIREMENT APPLIES AS WRITTEN TO EACH NTCB PARTITION THAT IS TRUSTED TO SEPARATE INFORMATION BASED ON SENSITIVITY LEVEL. EACH I/O DEVICE IN A COMPONENT, USED FOR COMMUNICATION WITH OTHER NETWORK COMPONENTS, IS ASSIGNED A DEVICE RANGE, CONSISTING OF A SET OF LABELS WITH A MAXIMUM AND MINIMUM. (A DEVICE RANGE USUALLY CONTAINS, BUT DOES NOT NECESSARILY CONTAIN, ALL POSSIBLE LABELS "BETWEEN" THE MAXIMUM AND MINIMUM, IN THE SENSE OF DOMINATING THE MINIMUM AND BEING DOMINATED BY THE MAXIMUM.)

THE NTCB ALWAYS PROVIDES AN ACCURATE LABEL FOR INFORMATION EXPORTED THROUGH DEVICES. INFORMATION EXPORTED OR IMPORTED USING A SINGLE-LEVEL DEVICE IS LABELLED IMPLICITLY BY THE SENSITIVITY LEVEL OF THE DEVICE. INFORMATION EXPORTED FROM ONE MULTILEVEL DEVICE AND IMPORTED AT ANOTHER MUST BE LABELLED THROUGH AN AGREED-UPON PROTOCOL, UNLESS IT IS LABELLED IMPLICITLY BY USING A COMMUNICATION LINK THAT ALWAYS CARRIES A SINGLE LEVEL.

INFORMATION EXPORTED AT A GIVEN SENSITIVITY LEVEL CAN BE SENT ONLY TO AN IMPORTING DEVICE WHOSE DEVICE RANGE CONTAINS THAT LEVEL OR A HIGHER LEVEL. IF THE IMPORTING DEVICE RANGE DOES NOT CONTAIN THE GIVEN LEVEL, THE INFORMATION IS RELABELLED UPON RECEPTION AT A HIGHER LEVEL WITHIN THE IMPORTING DEVICE RANGE. RELABELLING SHOULD NOT OCCUR OTHERWISE.

+ Rationale

THE PURPOSE OF DEVICE LABELS IS TO REFLECT AND CONSTRAIN THE SENSITIVITY LEVELS OF INFORMATION AUTHORIZED FOR THE PHYSICAL ENVIRONMENT IN WHICH THE DEVICES ARE LOCATED.

THE INFORMATION TRANSFER RESTRICTIONS PERMIT ONE-WAY COMMUNICATION (I.E., NO ACKNOWLEDGEMENTS) FROM ONE DEVICE TO ANOTHER WHOSE RANGES HAVE NO LEVEL IN COMMON, AS LONG AS EACH LEVEL IN THE SENDING DEVICE RANGE IS DOMINATED BY SOME LEVEL IN THE RECEIVING DEVICE RANGE. IT IS NEVER PERMITTED TO SEND INFORMATION AT A GIVEN LEVEL TO A DEVICE WHOSE RANGE DOES NOT CONTAIN A DOMINATING LEVEL. (SEE APPENDIX C FOR SIMILAR INTERCONNECTION RULES FOR THE INTERCONNECTED AIS VIEW.)

3.2.1.4 Mandatory Access Control

+ Statement from DoD 5200.28-STD

The TCB shall enforce a mandatory access control policy over all RESOURCES (I.E., SUBJECTS, STORAGE OBJECTS, AND I/O DEVICES) THAT ARE DIRECTLY OR INDIRECTLY ACCESSIBLE BY SUBJECTS

EXTERNAL TO THE TCB. These subjects and objects shall be assigned sensitivity labels that are a combination of hierarchical classification levels and non-hierarchical categories, and the labels shall be used as the basis for mandatory access control decisions. The TCB shall be able to support two or more such sensitivity levels. (See the Mandatory Access Control interpretations.) The following requirements shall hold for all accesses between ALL SUBJECTS EXTERNAL TO THE TCB AND ALL OBJECTS DIRECTLY OR INDIRECTLY ACCESSIBLE BY THESE SUBJECTS. A subject can read an object only if the hierarchical classification in the subject's sensitivity level is greater than or equal to the hierarchical classification of the object's sensitivity level and the non-hierarchical categories in the subject's sensitivity level include all the non-hierarchical categories in the object's sensitivity level. A subject can write an object only if the hierarchical classification in the subject's sensitivity level is less than or equal to the hierarchical classification of the object's sensitivity level and the non-hierarchical categories in the subject's sensitivity level are included in the non-hierarchical categories in the object's sensitivity level. Identification and authentication data shall be used by the TCB to authenticate the user's identity and to ensure that the sensitivity level and authorization of subjects external to the TCB that may be created to act on behalf of the individual user are dominated by the clearance and authorization of that user.

+ Interpretation

Each partition of the NTCB exercises mandatory access control policy over all subjects and objects in its COMPONENT. In a network, the responsibility of an NTCB partition encompasses all mandatory access control functions in its component that would be required of a TCB in a stand-alone system. In particular, subjects and objects used for communication with other components are under the control of the NTCB partition. Mandatory access control includes secrecy and integrity control to the extent that the network sponsor has described in the overall network security policy.

Conceptual entities associated with communication between two components, such as sessions, connections and virtual circuits, may be thought of as having two ends, one in each component, where each end is represented by a local object. Communication is viewed as an operation that copies information from an object at one end of a communication path to an object at the other end. Transient data-carrying entities, such as datagrams and packets, exist either as information within other objects, or as a pair of objects, one at each end of the communication path.

The requirement for "two or more" sensitivity levels can be met by either secrecy or integrity levels. When there is a mandatory integrity policy, the stated requirements for reading and writing are generalized to: A subject can read an object only if the subject's sensitivity level dominates the object's sensitivity level, and a subject can write an object only if the object's sensitivity level dominates the subject's sensitivity level. Based on the integrity policy, the network sponsor shall define the dominance relation for the total label, for example, by combining secrecy and integrity lattices. -

+ Rationale

An NTCB partition can maintain access control only over subjects and objects in its component. AT LEVELS B2 AND ABOVE, THE NTCB PARTITION MUST MAINTAIN ACCESS CONTROL OVER ALL SUBJECTS AND OBJECTS IN ITS COMPONENT. Access by a subject in one component to information contained in an object in another component requires the creation of a subject in the remote component which acts as a surrogate for the first subject.

The mandatory access controls must be enforced at the interface of the reference monitor (viz. the mechanism that controls physical processing resources) for each NTCB partition. This mechanism creates the abstraction of subjects and objects which it controls. Some of these subjects outside the reference monitor, per se, may be designated to implement part of an NTCB partition's mandatory policy, e.g., by using the "trusted subjects" defined in the Bell-LaPadula model.

The prior requirements on exportation of labeled information to and from I/O devices ensure the consistency between the sensitivity labels of objects connected by a communication path. As noted in the introduction, the network architecture must recognize the linkage between the overall mandatory network security policy and the connection oriented abstraction. For example, individual data-carrying entities such as datagrams can have individual sensitivity labels that subject them to mandatory access control in each component. The abstraction of a single-level connection is realized and enforced implicitly by an architecture while a connection is realized by single-level subjects that necessarily employ only datagrams of the same level.

The fundamental trusted systems technology permits the DAC mechanism to be distributed, in contrast to the requirements for mandatory access control. For networks this separation of MAC and DAC mechanisms is the rule rather than

- See, for example, Grohn, M. J., A Model of a Pro-

tected Data Management System, ESD-TR-76-289, I. P.

Sharp Assoc. Ltd., June, 1976; and Denning, D .E., Lunt, T. F., Neumann, P. G., Schell, R. R., Heckman, M. and Shockley, W., Secure Distributed Data Views, Security Policy and Interpretation for a Class A1 Multilevel Secure Relational Database System, SRI International, November 1986.

the exception.

The set of total sensitivity labels used to represent all the sensitivity levels for the mandatory access control (combined data secrecy and data integrity) policy always forms a partially ordered set. Without loss of generality, this set of labels can always be extended to form a lattice, by including all the combinations of non-hierarchical categories. As for any lattice, a dominance relation is always defined for the total sensitivity labels. For administrative reasons it may be helpful to have a maximum level which dominates all others.

3.2.2 Accountability

3.2.2.1 Identification and Authentication

+ Statement from DoD 5200.28-STD

The TCB shall require users to identify themselves to it before beginning to perform any other actions that the TCB is expected to mediate. Furthermore, the TCB shall maintain authentication data that includes information for verifying the identify of individual users (e.g., passwords) as well as information for determining the clearance and authorizations of individual users. This data shall be used by the TCB to authenticate the user's identity and to ensure that the sensitivity level and authorization of subjects external to the TCB that may be created to act on behalf of the individual user are dominated by the clearance and authorization of that user. The TCB shall protect authentication data so that it cannot be accessed by any unauthorized user. The TCB shall be able to enforce individual accountability by providing the capability to uniquely identify each individual ADP system user. The TCB shall also provide the capability of associating this identity with all auditable actions taken by that individual.

+ Interpretation

The requirement for identification and authentication of users is the same for a network system as for an ADP system. The identification and authentication may be done by the component to which the user is directly connected or some other component, such as an identification and authentication server. Available techniques, such as those described in the Password Guideline=, are generally also applicable in the network context. However, in cases where the NTCB is expected to mediate actions of a host (or other network component) that is acting on behalf of a user or group of users, the NTCB may employ identification and

= Department of Defense Password Management Guide-
line, CSC-STD-002-85

authentication of the host (or other component) in lieu of identification and authentication of an individual user, so long as the component identifier implies a list of specific users uniquely associated with the identifier at the time of its use for authentication. This requirement does not apply to internal subjects.

Authentication information, including the identity of a user (once authenticated) may be passed from one component to another without reauthentication, so long as the NTCB protects (e.g., by encryption) the information from unauthorized disclosure and modification. This protection shall provide at least a similar level of assurance (or strength of mechanism) as pertains to the protection of the authentication mechanism and authentication data.

+ Rationale

The need for accountability is not changed in the context of a network system. The fact that the NTCB is partitioned over a set of components neither reduces the need nor imposes new requirements. That is, individual accountability is still the objective. Also, in the context of a network system at the (C2) level or higher "individual accountability" can be satisfied by identification of a host (or other component) so long as the requirement for traceability to individual users or a set of specific individual users with active subjects is satisfied. There is allowed to be an uncertainty in traceability because of elapsed time between changes in the group membership and the enforcement in the access control mechanisms. In addition, there is no need in a distributed processing system like a network to reauthenticate a user at each point in the network where a projection of a user (via the subject operating on behalf of the user) into another remote subject takes place.

The passing of identifiers and/or authentication information from one component to another is usually done in support to the implementation of the discretionary access control (DAC). This support relates directly to the DAC regarding access by a user to a storage object in a different NTCB partition than the one where the user was authenticated. Employing a forwarded identification implies additional reliance on the source and components along the path. If the authenticated identification is used as the basis of determining a sensitivity label for a subject, it must satisfy the Label Integrity criterion.

An authenticated identification may be forwarded between components and employed in some component to identify the sensitivity level associated with a subject created to act on behalf of the user so identified.

3.2.2.1.1 Trusted Path

+ Statement from DoD 5200.28-STD

THE TCB SHALL SUPPORT A TRUSTED COMMUNICATION PATH BETWEEN ITSELF AND USER FOR INITIAL LOGIN AND AUTHENTICATION. COMMUNICATIONS VIA THIS PATH SHALL BE INITIATED EXCLUSIVELY BY A USER.

+ Interpretation

A TRUSTED PATH IS SUPPORTED BETWEEN A USER (I.E., HUMAN) AND THE NTCB PARTITION IN THE COMPONENT TO WHICH THE USER IS DIRECTLY CONNECTED.

+ Rationale

WHEN A USER LOGS INTO A REMOTE COMPONENT, THE USER ID IS TRANSMITTED SECURELY BETWEEN THE LOCAL AND REMOTE NTCB PARTITIONS IN ACCORDANCE WITH THE REQUIREMENTS IN IDENTIFICATION AND AUTHENTICATION.

TRUSTED PATH IS NECESSARY IN ORDER TO ASSURE THAT THE USER IS COMMUNICATING WITH THE NTCB AND ONLY THE NTCB WHEN SECURITY RELEVANT ACTIVITIES ARE TAKING PLACE (E.G., AUTHENTICATE USER, SET CURRENT SESSION SENSITIVITY LEVEL). HOWEVER, TRUSTED PATH DOES NOT ADDRESS COMMUNICATIONS WITHIN THE NTCB, ONLY COMMUNICATIONS BETWEEN THE USER AND THE NTCB. IF, THEREFORE, A COMPONENT DOES NOT SUPPORT ANY DIRECT USER COMMUNICATION THEN THE COMPONENT NEED NOT CONTAIN MECHANISMS FOR ASSURING DIRECT NTCB TO USER COMMUNICATIONS.

THE REQUIREMENT FOR TRUSTED COMMUNICATION BETWEEN ONE NTCB PARTITION AND ANOTHER NCTB PARTITION IS ADDRESSED IN THE SYSTEM ARCHITECTURE SECTION. THESE REQUIREMENTS ARE

SEPARATE AND DISTINCT FROM THE USER TO NTCB COMMUNICATION REQUIREMENT OF A TRUSTED PATH. HOWEVER, IT IS EXPECTED THAT THIS TRUSTED COMMUNICATION BETWEEN ONE NTCB PARTITION AND ANOTHER NTCB PARTITION WILL BE USED IN CONJUNCTION WITH THE TRUSTED PATH TO IMPLEMENT TRUSTED COMMUNICATION BETWEEN THE USER AND THE REMOTE NTCB PARTITION.

3.2.2.2 Audit

+ Statement from DoD 5200.28-STD

The TCB shall be able to create, maintain, and protect from modification or unauthorized access or destruction an audit trail of accesses to the objects it protects. The audit data shall be protected by the TCB so that read access to it is limited to those who are authorized for audit data. The TCB shall be able to record the following types of events: use of identification and authentication mechanisms, introduction of objects into a user's address space (e.g., file open, program initiation), deletion of objects, actions taken by computer operators and system administrators and/or system security officers, and other security relevant events. The TCB shall also be able to audit any override of human-readable output markings. For each recorded event, the audit record shall identify: date and time of the event, user, type of event, and success or failure of the event. For identification/authentication events the origin of request (e.g., terminal ID) shall be included in the audit record. For events that introduce an object into a user's address space and for object deletion events the audit record shall include the name of the object and the object's sensitivity level. The ADP system administrator shall be able to selectively audit the actions of any one or more users based on individual identify and/or object sensitivity level. THE TCB SHALL BE ABLE TO AUDIT THE IDENTIFIED EVENTS THAT MAY BE USED IN THE EXPLOITATION OF COVERT STORAGE CHANNELS.

+ Interpretation

This criterion applies as stated. The sponsor must select which events are auditable. If any such events are not distinguishable by the NTCB alone (for example those identified in Part II), the audit mechanism shall provide an interface, which an authorized subject can invoke with parameters sufficient to produce an audit record. These audit records shall be distinguishable from those provided by the NTCB. In the context of a network system, "other security relevant events" (depending on network system architecture and network security policy) might be as follows:

1. Identification of each access event (e.g., establishing a connection or a connectionless association between processes in two hosts of the network) and its principal parameters (e.g., host identifiers of the two hosts involved in the access event and user identifier or host identifier of the user or host that is requesting the access event)
2. Identification of the starting and ending times of each access event using local time or global synchronized time
3. Identification of security-relevant exceptional conditions (e.g., potential violation of data integrity, such as misrouted datagrams) detected during the transactions between two hosts
4. Utilization of cryptographic variables
5. Changing the configuration of the network (e.g., a component leaving the network and rejoining)

In addition, identification information should be included in appropriate audit trail records, as necessary, to allow association of all related (e.g., involving the same network event) audit trail records (e.g., at different hosts) with each other. Furthermore, a component of the network system may provide the required audit capability (e.g., storage, retrieval, reduction, analysis) for other components that do not internally store audit data but transmit the audit data to some designated collection component. Provisions shall be made to control the loss of audit data due to unavailability of resources.

In the context of a network system, the "user's address space" is extended, for object introduction and deletion events, to include address spaces being employed on behalf of a remote user (or host). However, the focus remains on users in contrast to internal subjects as discussed in the DAC criterion. In addition, audit information must be stored in machine-readable form.

THE CAPABILITY MUST EXIST TO AUDIT THE IDENTIFIED EVENTS THAT MAY BE USED IN THE EXPLOITATION OF COVERT STORAGE CHANNELS. TO ACCOMPLISH THIS, EACH NTCB PARTITION MUST BE ABLE TO AUDIT THOSE EVENTS LOCALLY THAT MAY LEAD TO THE EXPLOITATION OF A COVERT STORAGE CHANNEL WHICH EXIST BECAUSE OF THE NETWORK.

+ Rationale

For remote users, the network identifiers (e.g., internet address) can be used as identifiers of groups of indivi-

dual users (e.g., "all users at Host A") to eliminate the maintenance that would be required if individual identification of remote users was employed. In this class (C2), however, it must be possible to identify (immediately or at some later time) the individuals represented by a group identifier. In all other respects, the interpretation is a straightforward extension of the criterion into the context of a network system. IDENTIFICATION OF COVERT CHANNEL EVENTS IS ADDRESSED IN THE COVERT CHANNEL ANALYSIS SECTION.

3.2.3 Assurance

— — — — —

3.2.3.1 Operational Assurance

3.2.3.1.1 System Architecture

+ Statement from DoD 5200.28-STD

THE TCB SHALL MAINTAIN A DOMAIN FOR ITS OWN EXECUTION THAT PROTECTS IT FROM EXTERNAL INTERFERENCE OR TAMPERING (E.G., BY MODIFICATION OF ITS CODE OR DATA STRUCTURES). THE TCB SHALL MAINTAIN PROCESS ISOLATION THROUGH THE PROVISION OF DISTINCT ADDRESS SPACES UNDER ITS CONTROL. THE TCB SHALL BE INTERNALLY STRUCTURED INTO WELL-DEFINED LARGELY INDEPENDENT MODULES. IT SHALL MAKE EFFECTIVE USE OF AVAILABLE HARDWARE TO SEPARATE THOSE ELEMENTS THAT ARE PROTECTION-CRITICAL FROM THOSE THAT ARE NOT. THE TCB MODULES SHALL BE DESIGNED SUCH THAT THE PRINCIPLE OF LEAST PRIVILEGE IS ENFORCED. FEATURES IN HARDWARE, SUCH AS SEGMENTATION, SHALL BE USED TO SUPPORT LOGICALLY DISTINCT STORAGE OBJECTS WITH SEPARATE ATTRIBUTES (NAMELY: READABLE, WRITABLE). THE USER INTERFACE TO THE TCB SHALL BE COMPLETELY DEFINED AND ALL ELEMENTS OF THE TCB IDENTIFIED.

+ Interpretation

The system architecture criterion must be met individually by all NTCB partitions. Implementation of the requirement that the NTCB maintain a domain for its own execution is achieved by having each NTCB partition maintain a domain for its own execution. Since each component is itself a distinct domain in the overall network system, this also satisfies the requirement for process isolation through distinct address spaces in the special case where a component has only a single subject.

THE NTCB MUST BE INTERNALLY STRUCTURED INTO WELL-DEFINED LARGELY INDEPENDENT MODULES AND MEET THE HARDWARE REQUIREMENTS. THIS IS SATISFIED BY HAVING EACH NTCB PARTI-

TION SO STRUCTURED. THE NTCB CONTROLS ALL NETWORK RESOURCES. THESE RESOURCES are the union of the sets of resources over which the NTCB partitions have control. Code and data structures belonging to the NTCB, transferred among NTCB subjects (i.e., subjects outside the reference monitor but inside the NTCB) belonging to different NTCB partitions, must be protected against external interference or tampering. For example, a cryptographic checksum or physical means may be employed to protect user authentication data exchanged between NTCB partitions.

EACH NTCB PARTITION MUST ENFORCE THE PRINCIPLE OF LEAST PRIVILEGE WITHIN ITS COMPONENT. ADDITIONALLY, THE NTCB MUST BE STRUCTURED SO THAT THE PRINCIPLE OF LEAST PRIVILEGE IS ENFORCED IN THE SYSTEM AS A WHOLE.

Each NTCB partition provides isolation of resources (within its component) in accord with the network system architecture and security policy so that "supporting elements" (e.g., DAC and user identification) for the security mechanisms of the network system are strengthened compared to C2, from an assurance point of view, through the provision of distinct address spaces under control of the NTCB.

As discussed in the Discretionary Access Control section, the DAC mechanism of a NTCB partition may be implemented at the interface of the reference monitor or may be distributed in subjects that are part of the NTCB in the same or different component. When distributed in NTCB subjects (i.e., when outside the reference monitor), the assurance requirements for the design and implementation of the DAC shall be those of class C2 for all networks of class C2 or above.

+ Rationale

THE REQUIREMENT THAT THE NTCB BE STRUCTURED INTO MODULES AND MEET THE HARDWARE REQUIREMENTS APPLIES WITHIN THE NTCB PARTITIONS IN THE VARIOUS COMPONENTS.

THE PRINCIPLE OF LEAST PRIVILEGE REQUIRES THAT EACH USER OR OTHER INDIVIDUAL WITH ACCESS TO THE SYSTEM BE GIVEN ONLY THOSE RESOURCES AND AUTHORIZATIONS REQUIRED FOR THE PERFORMANCE OF THIS JOB. IN ORDER TO ENFORCE THIS PRINCIPLE IN THE SYSTEM IT MUST BE ENFORCED IN EVERY NTCB PARTITION THAT SUPPORTS USERS OR OTHER INDIVIDUALS. FOR EXAMPLE, PROHIBITING ACCESS BY ADMINISTRATORS TO OBJECTS OUTSIDE THE NTCB PARTITION (E.G., GAMES) LESSENS THE OPPORTUNITY OF DAMAGE BY A TROJAN HORSE.

The requirement for the protection of communications between NTCB partitions is specifically directed to subjects that are part of the NTCB partitions. Any requirements for

such protection for the subjects that are outside the NTCB partitions are addressed in response to the integrity requirements of the security policy.

The provision of distinct address spaces under the control of the NTCB provides the ability to separate subjects according to sensitivity level. This requirement is introduced at B1 since it is an absolute necessity in order to implement mandatory access controls.

3.2.3.1.2 System Integrity

+ Statement from DoD 5200.28-STD

Hardware and/or software features shall be provided that can be used to periodically validate the correct operation of the on-site hardware and firmware elements of the TCB.

+ Interpretation

Implementation of the requirement is partly achieved by having hardware and/or software features that can be used to periodically validate the correct operation of the hardware and firmware elements of each component's NTCB partition. Features shall also be provided to validate the identity and correct operation of a component prior to its incorporation in the network system and throughout system operation. For example, a protocol could be designed that enables the components of the partitioned NTCB to exchange messages periodically and validate each other's correct response. The protocol shall be able to determine the remote entity's ability to respond. NTCB partitions shall provide the capability to report to network administrative personnel the failures detected in other NTCB partitions.

Intercomponent protocols implemented within a NTCB shall be designed in such a way as to provide correct operation in the case of failures of network communications or individual components. The allocation of mandatory and discretionary access control policy in a network may require communication between trusted subjects that are part of the NTCB partitions in different components. This communication is normally implemented with a protocol between the subjects as peer entities. Incorrect access within a component shall not result from failure of an NTCB partition to communicate with other components.

+ Rationale

The first paragraph of the interpretation is a straightforward extension of the requirement into the context of a network system and partitioned NTCB as defined for these network criteria.

NTCB protocols should be robust enough so that they permit the system to operate correctly in the case of localized failure. The purpose of this protection is to preserve the integrity of the NTCB itself. It is not unusual for one or more components in a network to be inoperative at any time, so it is important to minimize the effects of such failures on the rest of the network. Additional integrity and denial of service issues are addressed in Part II.

3.2.3.1.3 Covert Channel Analysis

+ Statement from DoD 5200.28-STD

THE SYSTEM DEVELOPER SHALL CONDUCT A THOROUGH SEARCH FOR COVERT STORAGE CHANNELS AND MAKE A DETERMINATION (EITHER BY ACTUAL MEASUREMENT OR BY ENGINEERING ESTIMATION) OF THE MAXIMUM BANDWIDTH OF EACH IDENTIFIED CHANNEL. (SEE THE COVERT CHANNELS GUIDELINE SECTION.)

+ Interpretation

THE REQUIREMENT, INCLUDING THE TCSEC COVERT CHANNEL GUIDELINE, APPLIES AS WRITTEN. IN A NETWORK, THERE ARE ADDITIONAL INSTANCES OF COVERT CHANNELS ASSOCIATED WITH COMMUNICATION BETWEEN COMPONENTS.

+ Rationale

THE EXPLOITATION OF NETWORK PROTOCOL INFORMATION (E.G., HEADERS) CAN RESULT IN COVERT STORAGE CHANNELS. THE TOPIC HAS BEEN ADDRESSED IN THE LITERATURE.-

- See, for example, Girling, C. G., "Covert Channels in LAN's," IEEE Transactions on Software Engineering, Vol. SE-13, No. 2, February 1987; and Padlipsky, M. A., Snow, D. P., and Karger, P. A., Limitations of End-to-End Encryption in Secure Computer Networks, MITRE Technical Report, MTR-3592, Vol. I, May 1978 (ESD TR

3.2.3.1.4 Trusted Facility Management

+ Statement from DoD 5200.28-STD

THE TCB SHALL SUPPORT SEPARATE OPERATOR AND ADMINISTRATOR FUNCTIONS.

+ Interpretation

THIS REQUIREMENT APPLIES AS WRITTEN TO BOTH THE NETWORK AS A WHOLE AND TO INDIVIDUAL COMPONENTS WHICH SUPPORT SUCH PERSONNEL.

+ Rationale

IT IS RECOGNIZED THAT BASED ON THE ALLOCATED POLICY ELEMENTS SOME COMPONENTS MAY OPERATE WITH NO HUMAN INTER-FACE.

3.2.3.2 Life-Cycle Assurance

3.2.3.2.1 Security Testing

+ Statement from DoD 5200.28-STD

The security mechanisms of the ADP system shall be tested and found to work as claimed in the system documentation. A team of individuals who thoroughly understand the specific implementation of the TCB shall subject its design documentation, source code, and object code to thorough analysis and testing. Their objectives shall be: to uncover all design and implementation flaws that would permit a subject external to the TCB to read, change, or delete data normally denied under the mandatory or discretionary security policy enforced by the TCB; as well as to assure that no subject (without authorization to do so) is able to cause the TCB to enter a state such that it is unable to respond to communications initiated by other users. THE TCB SHALL BE FOUND RELATIVELY RESISTANT TO PENETRATION. All discovered flaws shall be removed or neutralized and the TCB retested to demonstrate that they have been eliminated and that new flaws have not been introduced. TESTING SHALL DEMONSTRATE THAT THE TCB IMPLEMENTATION IS CONSISTENT WITH THE DESCRIPTIVE TOP-LEVEL SPECIFICATION. (See the Security Testing Guidelines.)

+ Interpretation

Testing of a component will require a testbed that exercises the interfaces and protocols of the component including tests under exceptional conditions. The testing of a security mechanism of the network system for meeting this criterion shall be an integrated testing procedure involving all components containing an NTCB partition that implement the given mechanism. This integrated testing is additional to any individual component tests involved in the evaluation of the network system. The sponsor should identify the allowable set of configurations including the sizes of the networks. Analysis or testing procedures and tools

shall be available to test the limits of these configurations. A change in configuration within the allowable set of configurations does not require retesting.

The testing of each component will include the introduction of subjects external to the NTCB partition for the component that will attempt to read, change, or delete data normally denied. If the normal interface to the component does not provide a means to create the subjects needed to conduct such a test, then this portion of the testing shall use a special version of the untrusted software for the component that results in subjects that make such attempts. The results shall be saved for test analysis. Such special versions shall have an NTCB partition that is identical to that for the normal configuration of the component under evaluation.

The testing of the mandatory controls shall include tests to demonstrate that the labels for information imported and/or exported to/from the component accurately represent the labels maintained by the NTCB partition for the component for use as the basis for its mandatory access control decisions. The tests shall include each type of device, whether single-level or multilevel, supported by the component.

THE NTCB MUST BE FOUND RELATIVELY RESISTANT TO PENETRATION. THIS APPLIES TO THE NTCB AS A WHOLE, AND TO EACH NTCB PARTITION IN A COMPONENT OF THIS CLASS.

+ Rationale

The phrase "no subject (without authorization to do so) is able to cause the TCB to enter a state such that it is unable to respond to communications initiated by other users" relates to the security services (Part II of this TNI) for the Denial of Service problem, and to correctness of the protocol implementations.

Testing is an important method available in this evaluation division to gain any assurance that the security mechanisms perform their intended function. A major purpose of testing is to demonstrate the system's response to inputs to the NTCB partition from untrusted (and possibly malicious) subjects.

In contrast to general purpose systems that allow for the dynamic creation of new programs and the introductions of new processes (and hence new subjects) with user specified security properties, many network components have no method for introducing new programs and/or processes during their normal operation. Therefore, the programs necessary for the testing must be introduced as special versions of the software rather than as the result of normal inputs by

the test team. However, it must be insured that the NTCB partition used for such tests is identical to the one under evaluation.

Sensitivity labels serve a critical role in maintaining the security of the mandatory access controls in the network. Especially important to network security is the role of the labels for information communicated between components - explicit labels for multilevel devices and implicit labels for single-level devices. Therefore the testing for correct labels is highlighted.

THE REQUIREMENT FOR TESTING TO DEMONSTRATE CONSISTENCY BETWEEN THE NTCB IMPLEMENTATION AND THE DTLs IS A STRAIGHT-FORWARD EXTENSION OF THE TCSEC REQUIREMENT INTO THE CONTEXT OF A NETWORK SYSTEM.

3.2.3.2.2 Design Specification and Verification

+ Statement from DoD 5200.28-STD

A FORMAL model of the security policy supported by the TCB shall be maintained over the life cycle of the ADP system THAT IS PROVEN and demonstrated to be consistent with its axioms. A DESCRIPTIVE TOP-LEVEL SPECIFICATION (DTLS) OF THE TCB SHALL BE MAINTAINED THAT COMPLETELY AND ACCURATELY DESCRIBES THE TCB IN TERMS OF EXCEPTIONS, ERROR MESSAGES, AND EFFECTS. IT SHALL BE SHOWN TO BE AN ACCURATE DESCRIPTION OF THE TCB INTERFACE.

+ Interpretation

The overall network security policy expressed in this model will provide the basis for the mandatory access control policy exercised by the NTCB over subjects and storage objects in the entire network. The policy will also be the basis for the discretionary access control policy exercised by the NTCB to control access of named users to named objects. Data integrity requirements addressing the effects of unauthorized MSM need not be included in this model. The overall network policy must be decomposed into policy elements that are allocated to appropriate components and used as the basis for the security policy model for those components.

The level of abstraction of the model, and the set of subjects and objects that are explicitly represented in the model, will be affected by the NTCB partitioning. Subjects and objects must be represented explicitly in the model for the partition if there is some network component whose NTCB partition exercises access control over them. The model shall be structured so that the axioms and entities

applicable to individual network components are manifest. Global network policy elements that are allocated to components shall be represented by the model for that component.

THE REQUIREMENTS FOR A NETWORK DTLS ARE GIVEN IN THE DESIGN DOCUMENTATION SECTION.

+ Rationale

The treatment of the model depends to a great extent on the degree of integration of the communications service into a distributed system. In a closely coupled distributed system, one might use a model that closely resembles one appropriate for a stand-alone computer system.

In other cases, the model of each partition will be expected to show the role of the NTCB partition in each kind of component. It will most likely clarify the model, although not part of the model, to show access restrictions implied by the system design; for example, subjects representing protocol entities might have access only to objects containing data units at the same layer of protocol. The allocation of subjects and objects to different protocol layers is a protocol design choice which need not be reflected in the security policy model.

3.2.3.2.3 Configuration Management

+ Statement from DoD 5200.28-STD

DURING DEVELOPMENT AND MAINTENANCE OF THE TCB, A CONFIGURATION MANAGEMENT SYSTEM SHALL BE IN PLACE THAT MAINTAINS CONTROL OF CHANGES TO THE DESCRIPTIVE TOP-LEVEL SPECIFICATION, OTHER DESIGN DATA, IMPLEMENTATION DOCUMENTATION, SOURCE CODE, THE RUNNING VERSION OF THE OBJECT CODE, AND TEST FIXTURES AND DOCUMENTATION. THE CONFIGURATION MANAGEMENT SYSTEM SHALL ASSURE A CONSISTENT MAPPING AMONG ALL DOCUMENTATION AND CODE ASSOCIATED WITH THE CURRENT VERSION OF THE TCB. TOOLS SHALL BE PROVIDED FOR GENERATION OF A NEW VERSION OF THE TCB FROM SOURCE CODE. ALSO AVAILABLE SHALL BE TOOLS FOR COMPARING A NEWLY GENERATED VERSION WITH THE PREVIOUS TCB VERSION IN ORDER TO ASCERTAIN THAT ONLY THE INTENDED CHANGES HAVE BEEN MADE IN THE CODE THAT WILL ACTUALLY BE USED AS THE NEW VERSION OF THE TCB.

+ Interpretation

THE REQUIREMENT APPLIES AS WRITTEN, WITH THE FOLLOWING EXTENSIONS:

1. A CONFIGURATION MANAGEMENT SYSTEM MUST BE IN PLACE

FOR EACH NTCB PARTITION.

2. A CONFIGURATION MANAGEMENT PLAN MUST EXIST FOR THE ENTIRE SYSTEM. IF THE CONFIGURATION MANAGEMENT SYSTEM IS MADE UP OF THE CONGLOMERATION OF THE CONFIGURATION MANAGEMENT SYSTEMS OF THE VARIOUS NTCB PARTITIONS, THEN THE CONFIGURATION MANAGEMENT PLAN MUST ADDRESS THE ISSUE OF HOW CONFIGURATION CONTROL IS APPLIED TO THE SYSTEM AS A WHOLE.

+ Rationale

EACH NTCB PARTITION MUST HAVE A CONFIGURATION MANAGEMENT SYSTEM IN PLACE, OR ELSE THERE WILL BE NO WAY FOR THE NTCB AS A WHOLE TO HAVE AN EFFECTIVE CONFIGURATION MANAGEMENT SYSTEM. THE OTHER EXTENSIONS ARE MERELY REFLECTIONS OF THE WAY THAT NETWORKS OPERATE IN PRACTICE.

3.2.4 Documentation.

3.2.4.1 Security Features User's Guide

+ Statement from DoD 5200.28-STD

A single summary, chapter, or manual in user documentation shall describe the protection mechanisms provided by the TCB, interpretations on their use, and how they interact with one another.

+ Interpretation

This user documentation describes user visible protection mechanisms at the global (network system) level and at the user interface of each component, and the interaction among these.

+ Rationale

The interpretation is an extension of the requirement into the context of a network system as defined for these network criteria. Documentation of protection mechanisms provided by individual components is required by the criteria for trusted computer systems that are applied as appropriate for the individual components.

3.2.4.2 Trusted Facility Manual

+ Statement from DoD 5200.28-STD

A manual addressed to the ADP system administrator shall present cautions about functions and privileges that should be controlled when running a secure facility. The procedures for examining and maintaining the audit files as well as the detailed audit record structure for each type of audit event shall be given. The manual shall describe the operator and administrator functions related to security, to include changing the security characteristics of a user. It shall provide interpretations on the consistent and effective use of the protection features of the system, how they interact, how to securely generate a new TCB, and facility procedures, warnings, and privileges that need to be controlled in order to operate the facility in a secure manner. THE TCB MODULES THAT CONTAIN THE REFERENCE VALIDATION MECHANISM SHALL BE IDENTIFIED. THE PROCEDURES FOR SECURE GENERATION OF A NEW TCB FROM SOURCE AFTER MODIFICATION OF ANY MODULES IN THE TCB SHALL BE DESCRIBED.

+ Interpretation

This manual shall contain specifications and procedures to assist the system administrator(s) maintain cognizance of the network configuration. These specifications and procedures shall address the following:

1. The hardware configuration of the network itself;
2. The implications of attaching new components to the network;
3. The case where certain components may periodically leave the network (e.g., by crashing, or by being disconnected) and then rejoin;
4. Network configuration aspects that can impact the security of the network system; (For example, the manual should describe for the network system administrator the interconnections among components that are consistent with the overall network system architecture.)
5. Loading or modifying NTCB software or firmware (e.g., down-line loading).
6. INCREMENTAL UPDATES; THAT IS, IT MUST EXPLICITLY INDICATE WHICH COMPONENTS OF THE NETWORK MAY CHANGE WITHOUT OTHERS ALSO CHANGING.

The physical and administrative environmental controls shall be specified. Any assumptions about security of a given network should be clearly stated (e.g., the fact that all communications links must be physically protected to a certain level).

THE COMPONENTS OF THE NETWORK THAT FORM THE NTCB MUST BE IDENTIFIED. FURTHERMORE, THE MODULES WITHIN AN NTCB PARTITION THAT CONTAIN THE REFERENCE VALIDATION MECHANISM (IF ANY) WITHIN THAT PARTITION MUST BE IDENTIFIED.

THE PROCEDURES FOR THE SECURE GENERATION OF A NEW VERSION (OR COPY) OF EACH NTCB PARTITION FROM SOURCE MUST BE DESCRIBED. THE PROCEDURES AND REQUIREMENTS FOR THE SECURE GENERATION OF THE NTCB NECESSITATED BY CHANGES IN THE NETWORK CONFIGURATION SHALL BE DESCRIBED.

+ Rationale

There may be multiple system administrators with diverse responsibilities. The technical security measures described by these criteria must be used in conjunction with other forms of security in order to achieve security of the network. Additional forms include administrative security, physical security, emanations security, etc.

Extension of this criterion to cover configuration aspects of the network is needed because, for example, proper interconnection of components is typically essential to achieve a correct realization of the network architecture.

As mentioned in the section on Label Integrity, cryptography is one common mechanism employed to protect communication circuits. Encryption transforms the representation of information so that it is unintelligible to unauthorized subjects. Reflecting this transformation, the sensitivity of the ciphertext is generally lower than the cleartext. If encryption methodologies are employed, they shall be approved by the National Security Agency (NSA).

The encryption algorithm and its implementation are outside the scope of these interpretations. This algorithm and implementation may be implemented in a separate device or may be a function of a subject in a component not dedicated to encryption. Without prejudice, either implementation packaging is referred to as an encryption mechanism herein.

THE REQUIREMENTS FOR DESCRIPTIONS OF NTCB GENERATION AND IDENTIFICATION OF MODULES AND COMPONENTS THAT FORM THE NTCB ARE STRAIGHTFORWARD EXTENSIONS OF THE TCSEC REQUIREMENTS INTO THE NETWORK CONTEXT. IN THOSE CASES WHERE THE VENDOR DOES NOT PROVIDE SOURCE CODE, AN ACCEPTABLE PROCEDURE SHALL BE TO REQUEST THE VENDOR TO PERFORM THE SECURE GENERATION.

3.2.4.3 Test Documentation

+ Statement from DoD 5200.28-STD

The system developer shall provide to the evaluators a document that describes the test plan, test procedures that show how the security mechanisms were tested, and results of the security mechanisms' functional testing. IT SHALL INCLUDE RESULTS OF TESTING THE EFFECTIVENESS OF THE METHODS USED TO REDUCE COVERT CHANNEL BANDWIDTHS.

+ Interpretation

The "system developer" is interpreted as "the network system sponsor". The description of the test plan should establish the context in which the testing was or should be conducted. The description should identify any additional test components that are not part of the system being evaluated. This includes a description of the test-relevant functions of such test components and a description of the interfacing of those test components to the system being evaluated. The description of the test plan should also demonstrate that the tests adequately cover the network security policy. The tests should include the features described in the System Architecture and the System Integrity sections. The tests should also include network configuration and sizing.

+ Rationale

The entity being evaluated may be a networking subsystem (see Appendix A) to which other components must be added to make a complete network system. In that case, this interpretation is extended to include contextual definition because, at evaluation time, it is not possible to validate the test plans without the description of the context for testing the networking subsystem.

THE BANDWIDTHS OF COVERT CHANNELS ARE USED TO DETERMINE THE SUITABILITY OF A NETWORK SYSTEM FOR A GIVEN ENVIRONMENT. THE EFFECTIVENESS OF THE METHODS USED TO REDUCE THESE BANDWIDTHS MUST THEREFORE BE ACCURATELY DETERMINED.

3.2.4.4 Design Documentation

+ Statement from DoD 5200.28-STD

Documentation shall be available that provides a description of the manufacturer's philosophy of protection and an explanation of how this philosophy is translated into the TCB. THE interfaces between THE TCB modules shall be described. A FORMAL description of the security policy model enforced by the TCB shall be available and an explanation provided to

show that it is sufficient to enforce the security policy. The specific TCB protection mechanisms shall be identified and an explanation given to show that they satisfy the model. THE DESCRIPTIVE TOP-LEVEL SPECIFICATION (DTLS) SHALL BE SHOWN TO BE AN ACCURATE DESCRIPTION OF THE TCB INTERFACE. DOCUMENTATION SHALL DESCRIBE HOW THE TCB IMPLEMENTS THE REFERENCE MONITOR CONCEPT AND GIVE AN EXPLANATION WHY IT IS TAMPER RESISTANT, CANNOT BE BYPASSED, AND IS CORRECTLY IMPLEMENTED. DOCUMENTATION SHALL DESCRIBE HOW THE TCB IS STRUCTURED TO FACILITATE TESTING AND TO ENFORCE LEAST PRIVILEGE. THIS DOCUMENTATION SHALL ALSO PRESENT THE RESULTS OF THE COVERT CHANNEL ANALYSIS AND THE TRADEOFFS INVOLVED IN RESTRICTING THE CHANNELS. ALL AUDITABLE EVENTS THAT MAY BE USED IN THE EXPLOITATION OF KNOWN COVERT STORAGE CHANNELS SHALL BE IDENTIFIED. THE BANDWIDTHS OF KNOWN COVERT STORAGE CHANNELS, THE USE OF WHICH IS NOT DETECTABLE BY THE AUDITING MECHANISMS, SHALL BE PROVIDED. (SEE THE COVERT CHANNEL GUIDELINE SECTION.)

+ Interpretation

Explanation of how the sponsor's philosophy of protection is translated into the NTCB shall include a description of how the NTCB is partitioned. The security policy also shall be stated. The description of the interfaces between the NTCB modules shall include the interface(s) between NTCB partitions and modules within the partitions if the modules exist. The sponsor shall describe the security architecture and design, including the allocation of security requirements among components.

THE DOCUMENTATION INCLUDES BOTH A SYSTEM DESCRIPTION AND A SET OF COMPONENT DTLS'S. THE SYSTEM DESCRIPTION ADDRESSES THE NETWORK SECURITY ARCHITECTURE AND DESIGN BY SPECIFYING THE TYPES OF COMPONENTS IN THE NETWORK, WHICH ONES ARE TRUSTED, AND IN WHAT WAY THEY MUST COOPERATE TO SUPPORT NETWORK SECURITY OBJECTIVES. A COMPONENT DTLS SHALL BE PROVIDED FOR EACH TRUSTED NETWORK COMPONENT, I.E., EACH COMPONENT CONTAINING AN NTCB PARTITION. EACH COMPONENT DTLS SHALL DESCRIBE THE INTERFACE TO THE NTCB PARTITION OF ITS COMPONENT. APPENDIX A ADDRESSES COMPONENT EVALUATION ISSUES.

As stated in the introduction to Division B, the sponsor must demonstrate that the NTCB employs the reference monitor concept. The security policy model must be a model for a reference monitor.

The security policy model for each partition implementing a reference monitor shall fully represent the access control policy supported by the partition, including the discretionary and mandatory security policy for secrecy and/or integrity. For the mandatory policy the single dominance relation for sensitivity labels, including secrecy and/or integrity components, shall be precisely defined.

+ Rationale

The interpretation is a straightforward extension of the requirement into the context of a network system as defined for this network interpretation. Other documentation, such as description of components and description of operating environment(s) in which the networking subsystem or network system is designed to function, is required elsewhere, e.g., in the Trusted Facility Manual.

In order to be evaluated, a network must possess a coherent Network Security Architecture and Design. (Interconnection of components that do not adhere to such a single coherent Network Security Architecture is addressed in the Interconnection of Accredited AIS, Appendix C.) The Network Security Architecture must address the security-relevant policies, objectives, and protocols. The Network Security Design specifies the interfaces and services that must be incorporated into the network so that it can be evaluated as a trusted entity. There may be multiple designs that conform to the same architecture but are more or less incompatible and non-interoperable (except through the Interconnection Rules). Security related mechanisms requiring cooperation among components are specified in the design in terms of their visible interfaces; mechanisms having no visible interfaces are not specified in this document but are left as implementation decisions.

The Network Security Architecture and Design must be available from the network sponsor before evaluation of the network, or any component, can be undertaken. The Network Security Architecture and Design must be sufficiently complete, unambiguous, and free from obvious flaws to permit the construction or assembly of a trusted network based on the structure it specifies.

When a component is being designed or presented for evaluation, or when a network assembled from components is assembled or presented for evaluation, there must be a priori evidence that the Network security Architecture and Design are satisfied. That is, the components can be assembled into a network that conforms in every way with the Network Security Architecture and Design to produce a physical realization that is trusted to the extent that its evaluation indicates.

In order for a trusted network to be constructed from components that can be built independently, the Network Security Architecture and Design must completely and unambiguously define the security functionality of components as well as the interfaces between or among components. The Network Security Architecture and Design must be evaluated to determine that a network constructed to its specifica-

tions will in fact be trusted, that is, it will be evaluable under these interpretations.

The term "model" is used in several different ways in a network context, e.g., a "protocol reference model," a "formal network model," etc. Only the "security policy model" is addressed by this requirement and is specifically intended to model the interface, viz., "security perimeter," of the reference monitor and must meet all the requirements defined in the TCSEC. It must be shown that all parts of the TCB are a valid interpretation of the security policy model, i.e., that there is no change to the secure state except as represented by the model.

3.3 CLASS (B3): SECURITY DOMAINS

THE CLASS (B3) NTCB MUST SATISFY THE REFERENCE MONITOR REQUIREMENTS THAT IT MEDIATE ALL ACCESSES OF SUBJECTS TO OBJECTS, BE TAMPERPROOF, AND BE SMALL ENOUGH TO BE SUBJECTED TO ANALYSIS AND TESTS. TO THIS END, THE NTCB IS STRUCTURED TO EXCLUDE CODE NOT ESSENTIAL TO SECURITY POLICY ENFORCEMENT, WITH SIGNIFICANT SYSTEM ENGINEERING DURING NTCB DESIGN AND IMPLEMENTATION DIRECTED TOWARD MINIMIZING ITS COMPLEXITY. A SECURITY ADMINISTRATOR IS SUPPORTED, AUDIT MECHANISMS ARE EXPANDED TO SIGNAL SECURITY-RELEVANT EVENTS, AND SYSTEM RECOVERY PROCEDURES ARE REQUIRED. THE SYSTEM IS HIGHLY RESISTANT TO PENETRATION. THE FOLLOWING ARE MINIMAL REQUIREMENTS FOR SYSTEMS ASSIGNED A CLASS (B3) RATING:

3.3.1 Security Policy

+ Statement from DoD 5200.28-STD

Implied from the Introduction to the TCSEC.

+ Interpretation

The network sponsor shall describe the overall network security policy enforced by the NTCB. At a minimum, this policy shall include the discretionary and mandatory requirements applicable to this class. The policy may require data secrecy, or data integrity, or both. The policy is an access control policy having two primary components: mandatory and discretionary. The policy shall include a discretionary policy for protecting the informa-

tion being processed based on the authorizations of individuals, users, or groups of users. This access control policy statement shall describe the requirements on the network to prevent or detect "reading or destroying" sensitive information by unauthorized users or errors. The mandatory policy must define the set of distinct sensitivity levels that it supports. For the Class B1 or above the mandatory policy shall be based on the labels associated with the information that reflects its sensitivity with respect to secrecy and/or integrity, where applicable, and labels associated with users to reflect their authorization to access such information. Unauthorized users include both those that are not authorized to use the network at all (e.g., a user attempting to use a passive or active wire tap) or a legitimate user of the network who is not authorized to access a specific piece of information being protected.

Note that "users" does not include "operators," "system programmers," "technical control officers," "system security officers," and other system support personnel. They are distinct from users and are subject to the Trusted Facility Manual and the System Architecture requirements. Such individuals may change the system parameters of the network system, for example, by defining membership of a group. These individuals may also have the separate role of users.

SECRECY POLICY: The network sponsor shall define the form of the discretionary and mandatory secrecy policy that is enforced in the network to prevent unauthorized users from reading the sensitive information entrusted to the network.

DATA INTEGRITY POLICY: The network sponsor shall define the discretionary and mandatory integrity policy to prevent unauthorized users from modifying, viz., writing, sensitive information. The definition of data integrity presented by the network sponsor refers to the requirement that the information has not been subjected to unauthorized modification in the network. The mandatory integrity policy enforced by the NTCB cannot, in general, prevent modification while information is being transmitted between components. However, an integrity sensitivity label may reflect the confidence that the information has not been subjected to transmission errors because of the protection afforded during transmission. This requirement is distinct from the requirement for label integrity.

+ Rationale

The word "sponsor" is used in place of alternatives

(such as "vendor," "architect," "manufacturer," and "developer") because the alternatives indicate people who may not be available, involved, or relevant at the time that a network system is proposed for evaluation.

A trusted network is able to control both the reading and writing of shared sensitive information. Control of writing is used to protect against destruction of information. A network normally is expected to have policy requirements to protect both the secrecy and integrity of the information entrusted to it. In a network the integrity is frequently as important or more important than the secrecy requirements. Therefore the secrecy and/or integrity policy to be enforced by the network must be stated for each network regardless of its evaluation class. The assurance that the policy is faithfully enforced is reflected in the evaluation class of the network.

This control over modification is typically used to protect information so that it may be relied upon and to control the potential harm that would result if the information were corrupted. The overall network policy requirements for integrity includes the protection for data both while being processed in a component and while being transmitted in the network. The access control policy enforced by the NTCB relates to the access of subjects to objects within each component. Communications integrity addressed within Part II relates to information while being transmitted.

The mandatory integrity policy (at class B1 and above) in some architectures may be useful in supporting the linkage between the connection oriented abstraction introduced in the Introduction and the individual components of the network. For example, in a key distribution center for end-to-end encryption, a distinct integrity category may be assigned to isolate the key generation code and data from possible modification by other supporting processes in the same component, such as operator interfaces and audit.

The mandatory integrity policy for some architecture may define an integrity sensitivity label that reflects the specific requirements for ensuring that information has not been subject to random errors in excess of a stated limit nor to unauthorized message stream modification (MSM) -. The specific metric associated with an integrity sensitivity label will generally reflect the intended applications of the network.

3.3.1.1 Discretionary Access Control

+ Statement from DoD 5200.28-STD

The TCB shall define and control access between named users and named objects (e.g., files and programs) in the ADP system. The enforcement mechanism (e.g., ACCESS CONTROL LISTS) shall allow users to specify and control sharing of those OBJECTS and shall provide controls to limit propagation of access rights. The discretionary access control mechanism shall, either by explicit user action or by default, provide that objects are protected from unauthorized access. These access controls shall be capable of SPECIFYING, FOR EACH

- See Voydock, Victor L. and Stephen T. Kent, "Security Mechanisms in High-Level Network Protocols," Computing Surveys, Vol. 15, No. 2, June 1983, pp 135-171.

NAMED OBJECT, A LIST OF NAMED INDIVIDUALS AND A LIST OF GROUPS OF NAMED INDIVIDUALS WITH THEIR RESPECTIVE MODES OF ACCESS TO THAT OBJECT. FURTHERMORE, FOR EACH SUCH NAMED OBJECT, IT SHALL BE POSSIBLE TO SPECIFY A LIST OF NAMED INDIVIDUALS AND A LIST OF GROUPS OF NAMED INDIVIDUALS FOR WHICH NO ACCESS TO THE OBJECT IS GIVEN. Access permission to an object by users not already possessing access permission shall only be assigned by authorized users.

+ Interpretation

The discretionary access control (DAC) mechanism(s) may be distributed over the partitioned NTCB in various ways. Some part, all, or none of the DAC may be implemented in a given component of the network system. In particular, components that support only internal subjects (i.e., that have no subjects acting as direct surrogates for users), such as a public network packet switch, might not implement the DAC mechanism(s) directly (e.g., they are unlikely to contain access control lists).

Identification of users by groups may be achieved in various ways in the networking environment. For example, the network identifiers (e.g., internet addresses) for various components (e.g., hosts, gateways) can be used as identifiers of groups of individual users (e.g., "all users at Host A," "all users of network Q") so long as the individuals involved in the group are implied by the group identifier. For example, Host A might employ a particular group-id, for which it maintains a list of explicit users in that group, in its network exchange with Host B, which accepts the group-id under the conditions of this interpretation.

For networks, individual hosts will impose need-to-know controls over their users on the basis of named individuals - much like (in fact, probably the same) controls used when

there is no network connection.

When group identifiers are acceptable for access control, the identifier of some other host may be employed, to eliminate the maintenance that would be required if individual identification of remote users was employed. In class C2 and higher, however, it must be possible from that audit record to identify (immediately or at some later time) exactly the individuals represented by a group identifier at the time of the use of that identifier. There is allowed to be an uncertainty because of elapsed time between changes in the group membership and the enforcement in the access control mechanisms.

The DAC mechanism of a NTCB partition may be implemented at the interface of the reference monitor or may be distributed in subjects that are part of the NTCB in the same or different component. The reference monitor manages all the physical resources of the system and from them creates the abstraction of subjects and objects that it controls. Some of these subjects and objects may be used to implement a part of the NTCB. When the DAC mechanism is distributed in such NTCB subjects (i.e., when outside the reference monitor), the assurance requirements (see the Assurance section) for the design and implementation of the DAC shall be those of class C2 for all networks of class C2 or above.

When integrity is included as part of the network discretionary security policy, the above interpretations shall be specifically applied to the controls over modification, viz, the write mode of access, within each component based on identified users or groups of users.

+ Rationale

In this class, the supporting elements of the overall DAC mechanism are required to isolate information (objects) that supports DAC so that it is subject to auditing requirements (see the System Architecture section). The use of network identifiers to identify groups of individual users could be implemented, for example, as an X.25 community of interest in the network protocol layer (layer 3). In all other respects, the supporting elements of the overall DAC mechanism are treated exactly as untrusted subjects are treated with respect to DAC in an ADP system, with the same result as noted in the interpretation.

A typical situation for DAC is that a surrogate process for a remote user will be created in some host for access to objects under the control of the NTCB partition within that host. The interpretation requires that a user identifier be assigned and maintained for each such process by the NTCB, so that access by a surrogate process is subject to essen-

tially the same discretionary controls as access by a process acting on behalf of a local user would be. However, within this interpretation a range of possible interpretations of the assigned user identification is permitted.

The most obvious situation would exist if a global database of network users were to be made permanently available on demand to every host, (i.e., a name server existed) so that all user identifications were globally meaningful.

It is also acceptable, however, for some NTCB partitions to maintain a database of locally-registered users for its own use. In such a case, one could choose to inhibit the creation of surrogate processes for locally unregistered users, or (if permitted by the local policy) alternatively, to permit the creation of surrogate processes with preselected user and group identifiers which, in effect, identify the process as executing on behalf of a member of a group of users on a particular remote host. The intent of the words concerning audit in the interpretation is to provide a minimally acceptable degree of auditability for cases such as the last described. What is required is that there be a capability, using the audit facilities provided by the network NTCB partitions involved, to determine who was logged in at the actual host of the group of remote users at the time the surrogate processing occurred.

Associating the proper user id with a surrogate process is the job of identification and authentication. This means that DAC is applied locally, with respect to the user id of the surrogate process. The transmission of the data back across the network to the user's host, and the creation of a copy of the data there, is not the business of DAC.

Components that support only internal subjects impact the implementation of the DAC by providing services by which information (e.g., a user-id) is made available to a component that makes a DAC decision. An example of the latter would be the case that a user at Host A attempts to access a file at Host B. The DAC decision might be (and usually would be) made at Host B on the basis of a user-id transmitted from Host A to Host B.

Unique user identification may be achieved by a variety of mechanisms, including (a) a requirement for unique identification and authentication on the host where access takes place; (b) recognition of fully qualified network addresses authenticated by another host and forwarded to the host where access takes place; or (c) administrative support of a network-wide unique personnel identifier that could be authenticated and forwarded by another host as in (b) above, or could be authenticated and forwarded by a dedicated network identification and authentication server. The protocols which implement (b) or (c) are subject to the System

Architecture requirements.

Network support for DAC might be handled in other ways than that described as "typical" above. In particular, some form of centralized access control is often proposed. An access control center may make all decisions for DAC, or it may share the burden with the hosts by controlling host-to-host connections, and leaving the hosts to decide on access to their objects by users at a limited set of remote hosts. In this case the access control center provides the linkage between the connection oriented abstraction (as discussed in the Introduction) and the overall network security policy for DAC. In all cases the enforcement of the decision must be provided by the host where the object resides.

There are two forms of distribution for the DAC mechanism: implementing portions of the DAC in separate components, and supporting the DAC in subjects contained within the NTCB partition in a component. Since "the ADP system" is understood to be "the computer network" as a whole, each network component is responsible for enforcing security in the mechanisms allocated to it to ensure secure implementation of the network security policy. For traditional host systems it is frequently easy to also enforce the DAC along with the MAC within the reference monitor, per se, although a few approaches, such as virtual machine monitors, support DAC outside this interface.

In contrast to the universally rigid structure of mandatory policies (see the Mandatory Access Control section), DAC policies tend to be very network and system specific, with features that reflect the natural use of the system. For networks it is common that individual hosts will impose controls over their local users on the basis of named individuals-much like the controls used when there is no network connection. However, it is difficult to manage in a centralized manner all the individuals using a large network. Therefore, users on other hosts are commonly grouped together so that the controls required by the network DAC policy are actually based on the identity of the hosts or other components. A gateway is an example of such a component.

The assurance requirements are at the very heart of the concept of a trusted system. It is the assurance that determines if a system or network is appropriate for a given environment, as reflected, for example, in the Environments Guideline-. In the case of monolithic systems that have DAC integral to the reference monitor, the assurance requirements for DAC are inseparable from those of the rest of the reference monitor. For networks there is typically a much clearer distinction due to distributed DAC. The rationale for making the distinction in this network interpretation is that if major trusted network components can be made signi-

ificantly easier to design and implement without reducing the ability to meet security policy, then trusted networks will be more easily available.

3.3.1.2 Object Reuse

+ Statement from DoD 5200.28-STD

All authorizations to the information contained within a storage object shall be revoked prior to initial assignment, allocation or reallocation to a subject from the TCB's pool of unused storage objects. No information, including encrypted representations of information, produced by a prior subject's actions is to be available to any subject that obtains access to an object that has been released back to the system.

+ Interpretation

The NTCB shall ensure that any storage objects that it controls (e.g., message buffers under the control of a NTCB

- Guidance for Applying the Department of Defense
Trusted Computer System Evaluation Criteria in Specific
Environments, CSC-STD-003-85.

partition in a component) contain no information for which a subject in that component is not authorized before granting access. This requirement must be enforced by each of the NTCB partitions.

+ Rationale

In a network system, storage objects of interest are things that the NTCB directly controls, such as message buffers in components. Each component of the network system must enforce the object reuse requirement with respect to the storage objects of interest as determined by the network security policy. For example, the DAC requirement in this division leads to the requirement here that message buffers be under the control of the NTCB partition. A buffer assigned to an internal subject may be reused at the discretion of that subject which is responsible for preserving the integrity of message streams. Such controlled objects may be implemented in physical resources, such as buffers, disk sectors, tape space, and main memory, in components such as network switches.

3.3.1.3 Labels

+ Statement from DoD 5200.28-STD

Sensitivity labels associated with each ADP system resource (e.g., subject, storage object, ROM) that is directly or indirectly accessible by subjects external to the TCB shall be maintained by the TCB. These labels shall be used as the basis for mandatory access control decisions. In order to import non-labeled data, the TCB shall request and receive from an authorized user the sensitivity level of the data, and all such actions shall be auditable by the TCB.

+ Interpretation

Non-labeled data imported under the control of the NTCB partition will be assigned a label constrained by the device labels of the single-level device used to import it. Labels may include secrecy and integrity- components in accordance with the overall network security policy described by the network sponsor. Whenever the term "label" is used throughout this interpretation, it is understood to include both components as applicable. Similarly, the terms "single-level" and "multilevel" are understood to be based on both the secrecy and integrity components of the policy. The mandatory integrity policy will typically have requirements, such as the probability of undetected message stream modification, that will be reflected in the label for the

- See, for example, Biba, K.J., "Integrity Consideration for Secure Computer Systems," ESD-TR-76-372, MTR-3153, The MITRE Corporation, Bedford, MA, April 1977.

data so protected. For example, when data is imported its integrity label may be assigned based on mechanisms, such as cryptography, used to provide the assurance required by the policy. The NTCB shall assure that such mechanism are protected from tampering and are always invoked when they are the basis for a label.

If the security policy includes an integrity policy, all activities that result in message-stream modification during transmission are regarded as unauthorized accesses in violation of the integrity policy. The NTCB shall have an automated capability for testing, detecting, and reporting those errors/corruptions that exceed specified network integrity policy requirements. Message-stream modification (MSM) countermeasures shall be identified. A technology of adequate strength shall be selected to resist MSM. If encryption methodologies are employed, they shall be approved by the National Security Agency.

All objects must be labeled within each component of the network that is trusted to maintain separation of multiple levels of information. The label associated with any objects associated with single-level components will be identical to the level of that component. Objects used to store network control information, and other network structures, such as routing tables, must be labeled to prevent unauthorized access and/or modification.

+ Rationale

The interpretation is an extension of the requirement into the context of a network system and partitioned NTCB as defined for these network interpretations. A single-level device may be regarded either as a subject or an object. A multilevel device is regarded as a trusted subject in which the security range of the subject is the minimum-maximum range of the data expected to be transmitted over the device.

The sensitivity labels for either secrecy or integrity or both may reflect non-hierarchical categories or hierarchical classification or both.

For a network it is necessary that this requirement be applied to all network system resources at the (B2) level and above.

The NTCB is responsible for implementing the network integrity policy, when one exists. The NTCB must enforce that policy by ensuring that information is accurately transmitted from source to destination (regardless of the number of intervening connecting points). The NTCB must be able to counter equipment failure, environmental disruptions, and actions by persons and processes not authorized to alter the data. Protocols that perform code or format conversion shall preserve the integrity of data and control information.

The probability of an undetected transmission error may be specified as part of the network security policy so that the acceptability of the network for its intended application may be determined. The specific metrics (e.g., probability of undetected modification) satisfied by the data can be reflected in the integrity sensitivity label associated with the data while it is processed within a component. It is recognized that different applications and operational environments (e.g., crisis as compared to logistic) will have different integrity requirements.

The network shall also have an automated capability of testing for, detecting, and reporting errors that exceed a threshold consistent with the operational mode requirements. The effectiveness of integrity countermeasures must be est-

blished with the same rigor as the other security-relevant properties such as secrecy.

Cryptography is often utilized as a basis to provide data integrity assurance. Mechanisms, such as Manipulation Detection Codes (MDC)-, may be used. The adequacy of the encryption or MDC algorithm, the correctness of the protocol logic, and the adequacy of implementation must be established in MSM countermeasures design.

3.3.1.3.1 Label Integrity

+ Statement from DoD 5200.28-STD

Sensitivity labels shall accurately represent sensitivity levels of the specific subjects or objects with which they are associated. When exported by the TCB, sensitivity labels shall accurately and unambiguously represent the internal labels and shall be associated with the information being exported.

+ Interpretation

The phrase "exported by the TCB" is understood to include transmission of information from an object in one component to an object in another component. Information transferred between NTCB partitions is addressed in the System Integrity Section. The form of internal and external (exported) sensitivity labels may differ, but the meaning shall be the same. The NTCB shall, in addition, ensure that correct association of sensitivity labels with the information being transported across the network is preserved.

- See Jueneman, R. R., "Electronic Document Authentication," IEEE Network Magazine, April 1987, pp 17-23.

As mentioned in the Trusted Facility Manual Section, encryption transforms the representation of information so that it is unintelligible to unauthorized subjects. Reflecting this transformation, the sensitivity level of the ciphertext is generally lower than the cleartext. It follows that cleartext and ciphertext are contained in different objects, each possessing its own label. The label of the cleartext must be preserved and associated with the ciphertext so that it can be restored when the cleartext is subsequently obtained by decrypting the ciphertext. If the cleartext is associated with a single-level device, the label of that cleartext may be implicit. The label may also

be implicit in the key.

When information is exported to an environment where it is subject to deliberate or accidental modification, the TCB shall support the means, such as cryptographic checksums, to assure the accuracy of the labels. When there is a mandatory integrity policy, the policy will define the meaning of integrity labels.

+ Rationale

Encryption algorithms and their implementation are outside the scope of these interpretations. Such algorithms may be implemented in a separate device or may be incorporated in a subject of a larger component. Without prejudice, either implementation packaging is referred to as an encryption mechanism herein. If encryption methodologies are employed in this regard, they shall be approved by the National Security Agency (NSA). The encryption process is part of the Network Trusted Computer Base partition in the components in which it is implemented.

The encryption mechanism is not necessarily a multilevel device or multilevel subject, as these terms are used in these criteria. The process of encryption is multilevel by definition. The cleartext and ciphertext interfaces carry information of different sensitivity. An encryption mechanism does not process data in the sense of performing logical or arithmetic operations on that data with the intent of producing new data. The cleartext and ciphertext interfaces on the encryption mechanism must be separately identified as being single-level or multilevel. If the interface is single-level, then the sensitivity of the data is established by a trusted individual and implicitly associated with the interface; the Exportation to Single-Level Devices criterion applies.

If the interface is multilevel, then the data must be labeled; the Exportation to Multilevel Devices criterion applies. The network architect is free to select an acceptable mechanism for associating a label with an object. With reference to encrypted objects, the following examples are possible:

1. Include a label field in the protocol definition of the object.
2. Implicitly associate the label with the object through the encryption key. That is, the encryption key uniquely identifies a sensitivity level. A single or private key must be protected at the level of the data that it encrypts.

3.3.1.3.2 Exportation of Labeled Information

+ Statement from DoD 5200.28-STD

The TCB shall designate each communication channel and I/O device as either single-level or multilevel. Any change in this designation shall be done manually and shall be auditable by the TCB. The TCB shall maintain and be able to audit any change in the sensitivity level or levels associated with a communications channel or I/O device.

+ Interpretation

Each communication channel and network component shall be designated as either single-level or multilevel. Any change in this designation shall be done with the cognizance and approval of the administrator or security officer in charge of the affected components and the administrator or security officer in charge of the NTCB. This change shall be auditable by the network. The NTCB shall maintain and be able to audit any change in the device labels associated with a single-level communication channel or the range associated with a multilevel communication channel or component. The NTCB shall also be able to audit any change in the set of sensitivity levels associated with the information which can be transmitted over a multilevel communication channel or component.

+ Rationale

Communication channels and components in a network are analogous to communication channels and I/O devices in stand-alone systems. They must be designated as either multilevel (i.e., able to distinguish and maintain separation among information of various sensitivity levels) or single-level. As in the TCSEC, single-level devices may only be attached to single-level channels.

The level or set of levels of information that can be sent to a component or over a communication channel shall only change with the knowledge and approval of the security officers (or system administrator, if there is no security officer) of the network, and of the affected components. This requirement ensures that no significant security-relevant changes are made without the approval of all affected parties.

3.3.1.3.2.1 Exportation to Multilevel Devices

+ Statement from DoD 5200.28-STD

When the TCB exports an object to a multilevel I/O device, the sensitivity label associated with that object shall also

be exported and shall reside on the same physical medium as the exported information and shall be in the same form (i.e., machine-readable or human-readable form). When the TCB exports or imports an object over a multilevel communications channel, the protocol used on that channel shall provide for the unambiguous pairing between the sensitivity labels and the associated information that is sent or received.

+ Interpretation

The components, including hosts, of a network shall be interconnected over "multilevel communication channels," multiple single-level communication channels, or both, whenever the information is to be protected at more than a single sensitivity level. The protocol for associating the sensitivity label and the exported information shall provide the only information needed to correctly associate a sensitivity level with the exported information transferred over the multilevel channel between the NTCB partitions in individual components. This protocol definition must specify the representation and semantics of the sensitivity labels (i.e., the machine-readable label must uniquely represent the sensitivity level).

The "unambiguous" association of the sensitivity level with the communicated information shall meet the same level of accuracy as that required for any other label within the NTCB, as specified in the criterion for Label Integrity. This may be provided by protected and highly reliable direct physical layer connections, or by traditional cryptographic link protection in which any errors during transmission can be readily detected, or by use of a separate channel. The range of information imported or exported must be constrained by the associated device labels.

+ Rationale

This protocol must specify the representation and semantics of the sensitivity labels. See the Mandatory Access Control Policies section in Appendix B. The multilevel device interface to (untrusted) subjects may be implemented either by the interface of the reference monitor, per se, or by a multilevel subject (e.g., a "trusted subject" as defined in the Bell-LaPadula Model) that provides the labels based on the internal labels of the NTCB partition.

The current state of the art limits the support for mandatory policy that is practical for secure networks. Reference monitor support to ensure the control over all the operations of each subject in the network must be completely provided within the single NTCB partition on which that subject interfaces to the NTCB. This means that the entire

portion of the "secure state" represented in the formal security policy model that may be changed by transitions invoked by this subject must be contained in the same component.

The secure state of an NTCB partition may be affected by events external to the component in which the NTCB partition resides (e.g., arrival of a message). The effect occurs asynchronously after being initiated by an event in another component or partition. For example, indeterminate delays may occur between the initiation of a message in one component, the arrival of the message in the NTCB partition in another component, and the corresponding change to the secure state of the second component. Since each component is executing concurrently, to do otherwise would require some sort of network-wide control to synchronize state transitions, such as a global network-wide clock for all processors; in general, such designs are not practical and probably not even desirable. Therefore, the interaction between NTCB partitions is restricted to just communications between pairs (at least logically) of devices-multilevel devices if the device(s) can send/receive data of more than a single level. For broadcast channels the pairs are the sender and intended receiver(s). However, if the broadcast channel carries multiple levels of information, additional mechanism (e.g., cryptochecksum maintained by the TCB) may be required to enforce separation and proper delivery.

A common representation for sensitivity labels is needed in the protocol used on that channel and understood by both the sender and receiver when two multilevel devices (in this case, in two different components) are interconnected. Each distinct sensitivity level of the overall network policy must be represented uniquely in these labels.

Within a monolithic TCB, the accuracy of the sensitivity labels is generally assured by simple techniques, e.g., very reliable connections over very short physical connections, such as on a single printed circuit board or over an internal bus. In many network environments there is a much higher probability of accidentally or maliciously introduced errors, and these must be protected against.

3.3.1.3.2.2 Exportation to Single-Level Devices

+ Statement from DoD 5200.28-STD

Single-level I/O devices and single-level communication channels are not required to maintain the sensitivity labels of the information they process. However, the TCB shall include a mechanism by which the TCB and an authorized user reliably communicate to designate the single sensitivity level of information imported or exported via single-level

communication channels or I/O devices.

+ Interpretation

Whenever one or both of two directly connected components is not trusted to maintain the separation of information of different sensitivity levels, or whenever the two directly connected components have only a single sensitivity level in common, the two components of the network shall communicate over a single-level channel. Single-level components and single-level communication channels are not required to maintain the sensitivity labels of the information they process. However, the NTCB shall include a reliable communication mechanism by which the NTCB and an authorized user (via a trusted path) or a subject within an NTCB partition can designate the single sensitivity level of information imported or exported via single-level communication channels or network components. The level of information communicated must equal the device level.

+ Rationale

Single-level communications channels and single-level components in networks are analogous to single level channels and I/O devices in stand-alone systems in that they are not trusted to maintain the separation of information of different sensitivity levels. The labels associated with data transmitted over those channels and by those components are therefore implicit; the NTCB associates labels with the data because of the channel or component, not because of an explicit part of the bit stream. Note that the sensitivity level of encrypted information is the level of the ciphertext rather than the original level(s) of the plaintext.

3.3.1.3.2.3 Labeling Human-Readable Output

+ Statement from DoD 5200.28-STD

The ADP system administrator shall be able to specify the printable label names associated with exported sensitivity labels. The TCB shall mark the beginning and end of all human-readable, paged, hardcopy output (e.g., line printer output) with human-readable sensitivity labels that properly represent the sensitivity of the output. The TCB shall, by default, mark the top and bottom of each page of human-readable, paged, hardcopy output (e.g., line printer output) with human-readable sensitivity labels that properly represent the sensitivity of the page. The TCB shall, by default and in an appropriate manner, mark other forms of human readable output (e.g., maps, graphics) with human-readable sensitivity labels that properly represent the sensitivity of the output. Any override of these markings defaults shall be auditable by the TCB.

1 The hierarchical classification component in human-readable sensitivity labels shall be equal to the greatest hierarchical classification of any of the information in the output that the labels refer to; the non-hierarchical category component shall include all of the non-hierarchical categories of the information in the output the labels refer to, but no other non-hierarchical categories.

+ Interpretation

This criterion imposes no requirement to a component that produces no human-readable output. For those that do produce human-readable output, each sensitivity level that is defined to the network shall have a uniform meaning across all components. The network administrator, in conjunction with any affected component administrator, shall be able to specify the human-readable label that is associated with each defined sensitivity level.

+ Rationale

The interpretation is a straightforward extension of the requirement into the context of a network system and partitioned NTCB as defined for these network interpretations.

3.3.1.3.3 Subject Sensitivity Labels

+ Statement from DoD 5200.28-STD

The TCB shall immediately notify a terminal user of each change in the sensitivity level associated with that user during an interactive session. A terminal user shall be able to query the TCB as desired for a display of the subject's complete sensitivity label.

+ Interpretation

An NTCB partition shall immediately notify a terminal user attached to its component of each change in the sensitivity level associated with that user.

+ Rationale

The local NTCB partition must ensure that the user understands the sensitivity level of information sent to and from a terminal. When a user has a surrogate process in another component, adjustments to its level may occur to maintain communication with the user. These changes may occur asynchronously. Such adjustments are necessitated by mandatory access control as applied to the objects involved in the communication path.

3.3.1.3.4 Device Labels

+ Statement from DoD 5200.28-STD

The TCB shall support the assignment of minimum and maximum sensitivity levels to all attached physical devices. These sensitivity levels shall be used by the TCB to enforce constraints imposed by the physical environments in which the devices are located.

+ Interpretation

This requirement applies as written to each NTCB partition that is trusted to separate information based on sensitivity level. Each I/O device in a component, used for communication with other network components, is assigned a device range, consisting of a set of labels with a maximum and minimum. (A device range usually contains, but does not necessarily contain, all possible labels "between" the maximum and minimum, in the sense of dominating the minimum and being dominated by the maximum.)

The NTCB always provides an accurate label for information exported through devices. Information exported or imported using a single-level device is labelled implicitly by the sensitivity level of the device. Information exported from one multilevel device and imported at another must be labelled through an agreed-upon protocol, unless it is labelled implicitly by using a communication link that always carries a single level.

Information exported at a given sensitivity level can be sent only to an importing device whose device range contains that level or a higher level. If the importing device range does not contain the given level, the information is relabelled upon reception at a higher level within the importing device range. Relabelling should not occur otherwise.

+ Rationale

The purpose of device labels is to reflect and constrain the sensitivity levels of information authorized for the physical environment in which the devices are located.

The information transfer restrictions permit one-way communication (i.e., no acknowledgements) from one device to another whose ranges have no level in common, as long as each level in the sending device range is dominated by some level in the receiving device range. It is never permitted to send information at a given level to a device whose range does not contain a dominating level. (See Appendix C for

similar interconnection rules for the interconnected AIS view.)

3.3.1.4 Mandatory Access Control

+ Statement from DoD 5200.28-STD

The TCB shall enforce a mandatory access control policy over all resources (i.e., subjects, storage objects, and I/O devices) that are directly or indirectly accessible by subjects external to the TCB. These subjects and objects shall be assigned sensitivity labels that are a combination of hierarchical classification levels and non-hierarchical categories, and the labels shall be used as the basis for mandatory access control decisions. The TCB shall be able to support two or more such sensitivity levels. (See the Mandatory Access Control interpretations.) The following requirements shall hold for all accesses between all subjects external to the TCB and all objects directly or indirectly accessible by these subjects. A subject can read an object only if the hierarchical classification in the subject's sensitivity level is greater than or equal to the hierarchical classification of the object's sensitivity level and the non-hierarchical categories in the subject's sensitivity level include all the non-hierarchical categories in the object's sensitivity level. A subject can write an object only if the hierarchical classification in the subject's sensitivity level is less than or equal to the hierarchical classification of the object's sensitivity level and the non-hierarchical categories in the subject's sensitivity level are included in the non-hierarchical categories in the object's sensitivity level. Identification and authentication data shall be used by the TCB to authenticate the user's identity and to ensure that the sensitivity level and authorization of subjects external to the TCB that may be created to act on behalf of the individual user are dominated by the clearance and authorization of that user.

+ Interpretation

Each partition of the NTCB exercises mandatory access control policy over all subjects and objects in its component. In a network, the responsibility of an NTCB partition encompasses all mandatory access control functions in its component that would be required of a TCB in a stand-alone system. In particular, subjects and objects used for communication with other components are under the control of the NTCB partition. Mandatory access control includes secrecy and integrity control to the extent that the network sponsor has described in the overall network security policy.

Conceptual entities associated with communication between two components, such as sessions, connections and virtual circuits, may be thought of as having two ends, one in each component, where each end is represented by a local object. Communication is viewed as an operation that copies information from an object at one end of a communication path to an object at the other end. Transient data-carrying entities, such as datagrams and packets, exist either as information within other objects, or as a pair of objects, one at each end of the communication path.

The requirement for "two or more" sensitivity levels can be met by either secrecy or integrity levels. When there is a mandatory integrity policy, the stated requirements for reading and writing are generalized to: A subject can read an object only if the subject's sensitivity level dominates the object's sensitivity level, and a subject can write an object only if the object's sensitivity level dominates the subject's sensitivity level. Based on the integrity policy, the network sponsor shall define the dominance relation for the total label, for example, by combining secrecy and integrity lattices. -

+ Rationale

An NTCB partition can maintain access control only over subjects and objects in its component. At levels B2 and above, the NTCB partition must maintain access control over all subjects and objects in its component. Access by a subject in one component to information contained in an object in another component requires the creation of a subject in the remote component which acts as a surrogate for the first subject.

The mandatory access controls must be enforced at the interface of the reference monitor (viz. the mechanism that controls physical processing resources) for each NTCB partition. This mechanism creates the abstraction of subjects and objects which it controls. Some of these subjects outside the reference monitor, per se, may be designated to implement part of an NTCB partition's mandatory policy, e.g., by using the "trusted subjects" defined in the Bell-LaPadula model.

The prior requirements on exportation of labeled information to and from I/O devices ensure the consistency between the sensitivity labels of objects connected by a communication path. As noted in the introduction, the network architecture must recognize the linkage between the overall mandatory network security policy and the connection oriented abstraction. For example, individual data-carrying entities such as datagrams can have individual sensitivity labels that subject them to mandatory access control in each component. The abstraction of a single-level connection is

realized and enforced implicitly by an architecture while a connection is realized by single-level subjects that necessarily employ only datagrams of the same level.

The fundamental trusted systems technology permits the DAC mechanism to be distributed, in contrast to the requirements for mandatory access control. For networks this separation of MAC and DAC mechanisms is the rule rather than

- See, for example, Grohn, M. J., A Model of a Protected Data Management System, ESD-TR-76-289, I. P. Sharp Assoc. Ltd., June, 1976; and Denning, D. E., Lunt, T. F., Neumann, P. G., Schell, R. R., Heckman, M. and Shockley, W., Secure Distributed Data Views, Security Policy and Interpretation for a Class A1 Multilevel Secure Relational Database System, SRI International, November 1986.

the exception.

The set of total sensitivity labels used to represent all the sensitivity levels for the mandatory access control (combined data secrecy and data integrity) policy always forms a partially ordered set. Without loss of generality, this set of labels can always be extended to form a lattice, by including all the combinations of non-hierarchical categories. As for any lattice, a dominance relation is always defined for the total sensitivity labels. For administrative reasons it may be helpful to have a maximum level which dominates all others.

3.3.2 Accountability

3.3.2.1 Identification and Authentication

+ Statement from DoD 5200.28-STD

The TCB shall require users to identify themselves to it before beginning to perform any other actions that the TCB is expected to mediate. Furthermore, the TCB shall maintain authentication data that includes information for verifying the identify of individual users (e.g., passwords) as well as information for determining the clearance and authorizations of individual users. This data shall be used by the TCB to authenticate the user's identity and to ensure that the sensitivity level and authorization of subjects external

to the TCB that may be created to act on behalf of the individual user are dominated by the clearance and authorization of that user. The TCB shall protect authentication data so that it cannot be accessed by any unauthorized user. The TCB shall be able to enforce individual accountability by providing the capability to uniquely identify each individual ADP system user. The TCB shall also provide the capability of associating this identity with all auditable actions taken by that individual.

+ Interpretation

The requirement for identification and authentication of users is the same for a network system as for an ADP system. The identification and authentication may be done by the component to which the user is directly connected or some other component, such as an identification and authentication server. Available techniques, such as those described in the Password Guideline⁼, are generally also applicable in the network context. However, in cases where the NTCB is expected to mediate actions of a host (or other network component) that is acting on behalf of a user or group of users, the NTCB may employ identification and

⁼ Department of Defense Password Management Guideline, CSC-STD-002-85

authentication of the host (or other component) in lieu of identification and authentication of an individual user, so long as the component identifier implies a list of specific users uniquely associated with the identifier at the time of its use for authentication. This requirement does not apply to internal subjects.

Authentication information, including the identity of a user (once authenticated) may be passed from one component to another without reauthentication, so long as the NTCB protects (e.g., by encryption) the information from unauthorized disclosure and modification. This protection shall provide at least a similar level of assurance (or strength of mechanism) as pertains to the protection of the authentication mechanism and authentication data.

+ Rationale

The need for accountability is not changed in the context of a network system. The fact that the NTCB is partitioned over a set of components neither reduces the need nor imposes new requirements. That is, individual accountability is still the objective. Also, in the context of a network system at the (C2) level or higher "individual accountability" can be satisfied by identification of a host (or

other component) so long as the requirement for traceability to individual users or a set of specific individual users with active subjects is satisfied. There is allowed to be an uncertainty in traceability because of elapsed time between changes in the group membership and the enforcement in the access control mechanisms. In addition, there is no need in a distributed processing system like a network to reauthenticate a user at each point in the network where a projection of a user (via the subject operating on behalf of the user) into another remote subject takes place.

The passing of identifiers and/or authentication information from one component to another is usually done in support to the implementation of the discretionary access control (DAC). This support relates directly to the DAC regarding access by a user to a storage object in a different NTCB partition than the one where the user was authenticated. Employing a forwarded identification implies additional reliance on the source and components along the path. If the authenticated identification is used as the basis of determining a sensitivity label for a subject, it must satisfy the Label Integrity criterion.

An authenticated identification may be forwarded between components and employed in some component to identify the sensitivity level associated with a subject created to act on behalf of the user so identified.

3.3.2.1.1 Trusted Path

+ Statement from DoD 5200.28-STD

The TCB shall support a trusted communication path between itself and USERS FOR USE WHEN A POSITIVE TCB-TO-USER CONNECTION IS REQUIRED (E.G., LOGIN, CHANGE SUBJECT SENSITIVITY LEVEL). Communications via this TRUSTED path shall be ACTIVATED exclusively by a USER OR THE TCB AND SHALL BE LOGICALLY AND UNMISTAKABLY DISTINGUISHABLE FROM OTHER PATHS.

+ Interpretation

A trusted path is supported between a user (i.e., human) and the NTCB partition in the component to which the user is directly connected.

+ Rationale

When a user logs into a remote component, the user id is transmitted securely between the local and remote NTCB partitions in accordance with the requirements in Identification and Authentication.

Trusted Path is necessary in order to assure that the user is communicating with the NTCB and only the NTCB when security relevant activities are taking place (e.g., authenticate user, set current session sensitivity level). However, Trusted Path does not address communications within the NTCB, only communications between the user and the NTCB. If, therefore, a component does not support any direct user communication then the component need not contain mechanisms for assuring direct NTCB to user communications.

The requirement for trusted communication between one NTCB partition and another NTCB partition is addressed in the System Architecture section. These requirements are separate and distinct from the user to NTCB communication requirement of a trusted path. However, it is expected that this trusted communication between one NTCB partition and another NTCB partition will be used in conjunction with the trusted path to implement trusted communication between the user and the remote NTCB partition.

3.3.2.2 Audit

+ Statement from DoD 5200.28-STD

The TCB shall be able to create, maintain, and protect from modification or unauthorized access or destruction an audit trail of accesses to the objects it protects. The audit data shall be protected by the TCB so that read access to it is limited to those who are authorized for audit data. The TCB shall be able to record the following types of events: use of identification and authentication mechanisms, introduction of objects into a user's address space (e.g., file open, program initiation), deletion of objects, actions taken by computer operators and system administrators and/or system security officers, and other security relevant events. The TCB shall also be able to audit any override of human-readable output markings. For each recorded event, the audit record shall identify: date and time of the event, user, type of event, and success or failure of the event. For identification/authentication events the origin of request (e.g., terminal ID) shall be included in the audit record. For events that introduce an object into a user's address space and for object deletion events the audit record shall include the name of the object and the object's sensitivity level. The ADP system administrator shall be able to selectively audit the actions of any one or more users based on individual identify and/or object sensitivity level. The TCB shall be able to audit the identified events that may be used in the exploitation of covert storage channels. THE TCB SHALL CONTAIN A MECHANISM THAT IS ABLE TO MONITOR THE OCCURRENCE OR ACCUMULATION OF SECURITY AUDITABLE EVENTS THAT MAY INDICATE AN IMMINENT VIOLATION OF SECURITY POLICY. THIS MECHANISM SHALL BE ABLE TO

IMMEDIATELY NOTIFY THE SECURITY ADMINISTRATOR WHEN THRESHOLDS ARE EXCEEDED AND, IF THE OCCURRENCE OR ACCUMULATION OF THESE SECURITY RELEVANT EVENTS CONTINUES, THE SYSTEM SHALL TAKE THE LEAST DISRUPTIVE ACTION TO TERMINATE THE EVENT.

+ Interpretation

This criterion applies as stated. The sponsor must select which events are auditable. If any such events are not distinguishable by the NTCB alone (for example those identified in Part II), the audit mechanism shall provide an interface, which an authorized subject can invoke with parameters sufficient to produce an audit record. These audit records shall be distinguishable from those provided by the NTCB. In the context of a network system, "other security relevant events" (depending on network system architecture and network security policy) might be as follows:

1. Identification of each access event (e.g., establishing a connection or a connectionless association between processes in two hosts of the network) and its principal parameters (e.g., host identifiers of the two hosts involved in the access event and user identifier or host identifier of the user or host that is requesting the access event)
2. Identification of the starting and ending times of each access event using local time or global synchronized time
3. Identification of security-relevant exceptional conditions (e.g., potential violation of data integrity, such as misrouted datagrams) detected during the transactions between two hosts
4. Utilization of cryptographic variables
5. Changing the configuration of the network (e.g., a component leaving the network and rejoining)

In addition, identification information should be included in appropriate audit trail records, as necessary, to allow association of all related (e.g., involving the same network event) audit trail records (e.g., at different hosts) with each other. Furthermore, a component of the network system may provide the required audit capability (e.g., storage, retrieval, reduction, analysis) for other components that do not internally store audit data but transmit the audit data to some designated collection component. Provisions shall be made to control the loss of audit data due to unavailability of resources.

In the context of a network system, the "user's address space" is extended, for object introduction and deletion events, to include address spaces being employed on behalf of a remote user (or host). However, the focus remains on users in contrast to internal subjects as discussed in the DAC criterion. In addition, audit information must be stored in machine-readable form.

The capability must exist to audit the identified events that may be used in the exploitation of covert storage channels. To accomplish this, each NTCB partition must be able to audit those events locally that may lead to the exploitation of a covert storage channel which exist because of the network.

THE SPONSOR SHALL IDENTIFY THE SPECIFIC AUDITABLE EVENTS THAT MAY INDICATE AN IMMINENT VIOLATION OF SECURITY POLICY. THE COMPONENT WHICH DETECTS THE OCCURRENCE OR ACCUMULATION OF SUCH EVENTS MUST BE ABLE TO NOTIFY AN APPROPRIATE ADMINISTRATOR WHEN THRESHOLDS ARE EXCEEDED, AND TO INITIATE ACTIONS WHICH WILL RESULT IN TERMINATION OF THE EVENT IF THE ACCUMULATION CONTINUES. FOR EXAMPLE, WHEN THE THRESHOLD OF UNSUCCESSFUL LOGIN ATTEMPTS WITHIN A PERIOD OF TIME IS EXCEEDED, LOGIN SHALL BE INHIBITED FOR A SPECIFIC TIME PERIOD.

+ Rationale

For remote users, the network identifiers (e.g., internet address) can be used as identifiers of groups of individual users (e.g., "all users at Host A") to eliminate the maintenance that would be required if individual identification of remote users was employed. In this class (C2), however, it must be possible to identify (immediately or at some later time) the individuals represented by a group identifier. In all other respects, the interpretation is a straightforward extension of the criterion into the context of a network system. Identification of covert channel events is addressed in the Covert Channel Analysis section.

BECAUSE OF CONCURRENCY AND SYNCHRONIZATION PROBLEMS, IT MAY NOT BE POSSIBLE TO DETECT IN REAL TIME THE ACCUMULATION OF SECURITY AUDITABLE EVENTS THAT ARE OCCURRING IN DIFFERENT NTCB PARTITIONS. HOWEVER, EACH NTCB PARTITION THAT HAS BEEN ALLOCATED AUDIT RESPONSIBILITY MUST HAVE THE CAPABILITY TO DETECT THE LOCAL ACCUMULATION OF EVENTS, TO NOTIFY THE PARTITION SECURITY ADMINISTRATOR AND/OR THE NETWORK SECURITY ADMINISTRATOR, AND TO INITIATE ACTIONS WHICH WILL RESULT IN TERMINATION OF THE EVENT LOCALLY.

3.3.3 Assurance

3.3.3.1 Operational Assurance

3.3.3.1.1 System Architecture

+ Statement from DoD 5200.28-STD

The TCB shall maintain a domain for its own execution that protects it from external interference or tampering (e.g., by modification of its code or data structures). The TCB shall maintain process isolation through the provision of distinct address spaces under its control. The TCB shall be internally structured into well-defined largely independent modules. It shall make effective use of available hardware to separate those elements that are protection-critical from those that are not. The TCB modules shall be designed such that the principle of least privilege is enforced. Features in hardware, such as segmentation, shall be used to support logically distinct storage objects with separate attributes (namely: readable, writable). The user interface to the TCB shall be completely defined and all elements of the TCB identified. THE TCB SHALL BE DESIGNED AND STRUCTURED TO USE A COMPLETE, CONCEPTUALLY SIMPLE PROTECTION MECHANISM WITH PRECISELY DEFINED SEMANTICS. THIS MECHANISM SHALL PLAY A CENTRAL ROLE IN ENFORCING THE INTERNAL STRUCTURING OF THE TCB AND THE SYSTEM. THE TCB SHALL INCORPORATE SIGNIFICANT USE OF LAYERING, ABSTRACTION AND DATA HIDING. SIGNIFICANT SYSTEM ENGINEERING SHALL BE DIRECTED TOWARD MINIMIZING THE COMPLEXITY OF THE TCB AND EXCLUDING FROM THE TCB MODULES THAT ARE NOT PROTECTION-CRITICAL.

+ Interpretation

The system architecture criterion must be met individually by all NTCB partitions. Implementation of the requirement that the NTCB maintain a domain for its own execution is achieved by having each NTCB partition maintain a domain for its own execution. Since each component is itself a distinct domain in the overall network system, this also satisfies the requirement for process isolation through distinct address spaces in the special case where a component has only a single subject.

The NTCB must be internally structured into well-defined largely independent modules and meet the hardware requirements. This is satisfied by having each NTCB partition so structured. The NTCB controls all network resources. These resources are the union of the sets of resources over which the NTCB partitions have control. Code and data structures belonging to the NTCB, transferred among NTCB subjects (i.e., subjects outside the reference monitor but inside the NTCB) belonging to different NTCB partitions, must be protected against external interference or tamper-

ing. For example, a cryptographic checksum or physical means may be employed to protect user authentication data exchanged between NTCB partitions.

Each NTCB partition must enforce the principle of least privilege within its component. Additionally, the NTCB must be structured so that the principle of least privilege is enforced in the system as a whole.

THE NTCB MUST BE DESIGNED AND STRUCTURED ACCORDING TO THE NETWORK SECURITY ARCHITECTURE TO USE A COMPLETE, CONCEPTUALLY SIMPLE PROTECTION MECHANISM. FURTHERMORE, EACH NTCB PARTITION MUST ALSO BE SO DESIGNED AND STRUCTURED.

SIGNIFICANT SYSTEM ENGINEERING SHOULD BE DIRECTED TOWARD MINIMIZING THE COMPLEXITY OF EACH NTCB PARTITION, AND OF THE NTCB. CARE SHALL BE TAKEN TO EXCLUDE MODULES (AND COMPONENTS) THAT ARE NOT PROTECTION-CRITICAL FROM THE NTCB.

IT IS RECOGNIZED THAT SOME MODULES AND/OR COMPONENTS MAY NEED TO BE INCLUDED IN THE NTCB AND MUST MEET THE NTCB REQUIREMENTS EVEN THOUGH THEY MAY NOT APPEAR TO BE DIRECTLY PROTECTION-CRITICAL. THE CORRECT OPERATION OF THESE MODULES/COMPONENTS IS NECESSARY FOR THE CORRECT OPERATION OF THE PROTECTION-CRITICAL MODULES AND COMPONENTS. HOWEVER, THE NUMBER AND SIZE OF THESE MODULES/COMPONENTS SHOULD BE KEPT TO A MINIMUM.

Each NTCB partition provides isolation of resources (within its component) in accord with the network system architecture and security policy so that "supporting elements" (e.g., DAC and user identification) for the security mechanisms of the network system are strengthened compared to C2, from an assurance point of view, through the provision of distinct address spaces under control of the NTCB.

As discussed in the Discretionary Access Control section, the DAC mechanism of a NTCB partition may be implemented at the interface of the reference monitor or may be distributed in subjects that are part of the NTCB in the same or different component. When distributed in NTCB subjects (i.e., when outside the reference monitor), the assurance requirements for the design and implementation of the DAC shall be those of class C2 for all networks of class C2 or above.

+ Rationale

The requirement that the NTCB be structured into modules and meet the hardware requirements applies within the NTCB partitions in the various components.

The principle of least privilege requires that each

user or other individual with access to the system be given only those resources and authorizations required for the performance of this job. In order to enforce this principle in the system it must be enforced in every NTCB partition that supports users or other individuals. For example, prohibiting access by administrators to objects outside the NTCB partition (e.g., games) lessens the opportunity of damage by a Trojan Horse.

The requirement for the protection of communications between NTCB partitions is specifically directed to subjects that are part of the NTCB partitions. Any requirements for such protection for the subjects that are outside the NTCB partitions are addressed in response to the integrity requirements of the security policy.

THERE ARE CERTAIN PARTS OF A NETWORK (MODULES AND/OR COMPONENTS) THAT MAY NOT APPEAR TO BE DIRECTLY PROTECTION-CRITICAL IN THAT THEY ARE NOT INVOLVED IN ACCESS CONTROL DECISIONS, DO NOT DIRECTLY AUDIT, AND ARE NOT INVOLVED IN THE IDENTIFICATION/AUTHENTICATION PROCESS. HOWEVER, THE SECURITY OF THE NETWORK MUST DEPEND ON THE CORRECT OPERATION OF THESE MODULES AND/OR COMPONENTS. AN EXAMPLE OF THIS IS A SINGLE LEVEL PACKET SWITCH. ALTHOUGH IT MAY NOT NORMALLY BE INVOLVED DIRECTLY IN ENFORCING THE DISCRETIONARY SECURITY POLICY, THIS SWITCH MAY BE TRUSTED NOT TO MIX DATA FROM DIFFERENT MESSAGE STREAMS. IF THE SWITCH DOES NOT OPERATE CORRECTLY, DATA COULD GET MIXED, AND UNAUTHORIZED ACCESS COULD RESULT. THEREFORE, THESE MODULES/COMPONENTS MUST BE INCLUDED IN THE NTCB AND MUST MEET THE NTCB REQUIREMENTS APPLICABLE TO THE POLICY ELEMENT(S) FOR WHICH THEY ARE RESPONSIBLE.

3.3.3.1.2 System Integrity

+ Statement from DoD 5200.28-STD

Hardware and/or software features shall be provided that can be used to periodically validate the correct operation of the on-site hardware and firmware elements of the TCB.

+ Interpretation

Implementation of the requirement is partly achieved by having hardware and/or software features that can be used to periodically validate the correct operation of the hardware and firmware elements of each component's NTCB partition. Features shall also be provided to validate the identity and correct operation of a component prior to its incorporation in the network system and throughout system operation. For example, a protocol could be designed that enables the components of the partitioned NTCB to exchange messages periodically and validate each other's correct response. The pro-

TOCOL shall be able to determine the remote entity's ability to respond. NTCB partitions shall provide the capability to report to network administrative personnel the failures detected in other NTCB partitions.

Intercomponent protocols implemented within a NTCB shall be designed in such a way as to provide correct operation in the case of failures of network communications or individual components. The allocation of mandatory and discretionary access control policy in a network may require communication between trusted subjects that are part of the NTCB partitions in different components. This communication is normally implemented with a protocol between the subjects as peer entities. Incorrect access within a component shall not result from failure of an NTCB partition to communicate with other components.

+ Rationale

The first paragraph of the interpretation is a straightforward extension of the requirement into the context of a network system and partitioned NTCB as defined for these network criteria.

NTCB protocols should be robust enough so that they permit the system to operate correctly in the case of localized failure. The purpose of this protection is to preserve the integrity of the NTCB itself. It is not unusual for one or more components in a network to be inoperative at any time, so it is important to minimize the effects of such failures on the rest of the network. Additional integrity and denial of service issues are addressed in Part II.

IT SHOULD BE CLEAR THAT SOME INTEGRITY AND DENIAL OF SERVICE FEATURES CAN RESIDE OUTSIDE THE NTCB. OTHERWISE ALL SOFTWARE IN A NETWORK WOULD BE IN THE NTCB. EVERY PIECE OF SOFTWARE THAT HAS AN OPPORTUNITY TO WRITE TO SOME DATA OR PROTOCOL FIELD IS "TRUSTED" TO PRESERVE INTEGRITY OR NOT CAUSE DENIAL OF SERVICE TO SOME EXTENT. FOR EXAMPLE, IT IS NECESSARY TO "TRUST" TELNET TO CORRECTLY TRANSLATE USER DATA, AND TO EVENTUALLY TRANSMIT PACKETS. FTP ALSO HAS TO BE "TRUSTED" TO NOT INAPPROPRIATELY MODIFY FILES, AND TO ATTEMPT TO COMPLETE THE FILE TRANSFER. THESE PROTOCOLS CAN BE DESIGNED, HOWEVER TO EXIST OUTSIDE THE NTCB (FROM A PROTECTION PERSPECTIVE). IT IS BENEFICIAL TO DO THIS TYPE OF SECURITY ENGINEERING SO THAT THE AMOUNT OF CODE THAT MUST BE TRUSTED TO NOT DISCLOSE DATA IS MINIMIZED. PUTTING EVERYTHING INSIDE THE NTCB CONTRADICTS THE REQUIREMENT TO PERFORM "SIGNIFICANT SYSTEM ENGINEERING ... DIRECTED TOWARD ... EXCLUDING FROM THE TCB MODULES THAT ARE NOT PROTECTION CRITICAL," WHICH REMOVES THE PRIMARY DIFFERENCE BETWEEN B2 AND B3. IF EVERYTHING HAS TO BE IN THE TCB TO ENSURE DATA INTEGRITY AND PROTECTION AGAINST DENIAL OF SERVICE, THERE WILL BE CONSIDERABLY LESS ASSURANCE THAT DISCLOSURE PROTEC-

TION IS MAXIMIZED.

3.3.3.1.3 Covert Channel Analysis

+ Statement from DoD 5200.28-STD

The system developer shall conduct a thorough search for COVERT CHANNELS and make a determination (either by actual measurement or by engineering estimation) of the maximum bandwidth of each identified channel. (See the Covert Channels Guideline section.)

+ Interpretation

The requirement, including the TCSEC Covert Channel Guideline, applies as written. In a network, there are additional instances of covert channels associated with communication between components.

+ Rationale

The exploitation of network protocol information (e.g., headers) can result in covert storage channels. EXPLOITATION OF FREQUENCY OF TRANSMISSION CAN RESULT IN COVERT TIMING CHANNELS. The topic has been addressed in the literature.-

3.3.3.1.4 Trusted Facility Management

+ Statement from DoD 5200.28-STD

The TCB shall support separate operator and administrator functions. THE FUNCTIONS PERFORMED IN THE ROLE OF A SECURITY ADMINISTRATOR SHALL BE IDENTIFIED. THE ADP SYSTEM ADMINISTRATIVE PERSONNEL SHALL ONLY BE ABLE TO PERFORM SECURITY ADMINISTRATOR FUNCTIONS AFTER TAKING A DISTINCT AUDITABLE ACTION TO ASSUME THE SECURITY ADMINISTRATOR ROLE ON THE ADP SYSTEM. NON-SECURITY FUNCTIONS THAT CAN BE PERFORMED IN THE SECURITY ADMINISTRATION ROLE SHALL BE LIMITED STRICTLY TO THOSE ESSENTIAL TO PERFORMING THE SECURITY ROLE EFFECTIVELY.

- See, for example, Girling, C. G., "Covert Channels in LAN's," IEEE Transactions on Software Engineering,

Vol. SE-13, No. 2, February 1987; and Padlipsky, M. A., Snow, D. P., and Karger, P. A., Limitations of End-to-

End Encryption in Secure Computer Networks, MITRE

Technical Report, MTR-3592, Vol. I, May 1978 (ESD TR 78-158, DTIC AD A059221).

+ Interpretation

This requirement applies as written to both the network as a whole and to individual components which support such personnel.

+ Rationale

It is recognized that based on the allocated policy elements some components may operate with no human interface.

3.3.3.1.5 Trusted Recovery

+ Statement from DoD 5200.28-STD

PROCEDURES AND/OR MECHANISMS SHALL BE PROVIDED TO ASSURE THAT, AFTER AN ADP SYSTEM FAILURE OR OTHER DISCONTINUITY, RECOVERY WITHOUT A PROTECTION COMPROMISE IS OBTAINED.

+ Interpretation

THE RECOVERY PROCESS MUST BE ACCOMPLISHED WITHOUT A PROTECTION COMPROMISE AFTER THE FAILURE OR OTHER DISCONTINUITY OF ANY NTCB PARTITION. IT MUST ALSO BE ACCOMPLISHED AFTER A FAILURE OF THE ENTIRE NTCB.

+ Rationale

THIS IS A STRAIGHT-FORWARD EXTENSION OF THE REQUIREMENT INTO THE NETWORK CONTEXT, AND TAKES INTO ACCOUNT THAT IT IS POSSIBLE FOR PARTS OF THE SYSTEM TO FAIL WHILE OTHER PARTS CONTINUE TO OPERATE NORMALLY. THIS MAY BE A SECURITY-RELEVANT EVENT; IF SO IT MUST BE AUDITED.

3.3.3.2 Life-Cycle Assurance

3.3.3.2.1 Security Testing

+ Statement from DoD 5200.28-STD

The security mechanisms of the ADP system shall be tested and found to work as claimed in the system documentation. A team of individuals who thoroughly understand the specific implementation of the TCB shall subject its design documen-

tation, source code, and object code to through analysis and testing. Their objectives shall be: to uncover all design and implementation flaws that would permit a subject external to the TCB to read, change, or delete data normally denied under the mandatory or discretionary security policy enforced by the TCB; as well as to assure that no subject (without authorization to do so) is able to cause the TCB to enter a state such that it is unable to respond to communications initiated by other users. The TCB shall be FOUND RESISTANT to penetration. All discovered flaws shall be removed or neutralized and the TCB retested to demonstrate that they have been eliminated and that new flaws have not been introduced. Testing shall demonstrate that the TCB implementation is consistent with the descriptive top-level specification. NO DESIGN FLAWS AND NO MORE THAN A FEW CORRECTABLE IMPLEMENTATION FLAWS MAY BE FOUND DURING TESTING AND THERE SHALL BE REASONABLE CONFIDENCE THAT FEW REMAIN. (See the Security Testing Guidelines.)

+ Interpretation

Testing of a component will require a testbed that exercises the interfaces and protocols of the component including tests under exceptional conditions. The testing of a security mechanism of the network system for meeting this criterion shall be an integrated testing procedure involving all components containing an NTCB partition that implement the given mechanism. This integrated testing is additional to any individual component tests involved in the evaluation of the network system. The sponsor should identify the allowable set of configurations including the sizes of the networks. Analysis or testing procedures and tools shall be available to test the limits of these configurations. A change in configuration within the allowable set of configurations does not require retesting.

The testing of each component will include the introduction of subjects external to the NTCB partition for the component that will attempt to read, change, or delete data normally denied. If the normal interface to the component does not provide a means to create the subjects needed to conduct such a test, then this portion of the testing shall use a special version of the untrusted software for the component that results in subjects that make such attempts. The results shall be saved for test analysis. Such special versions shall have an NTCB partition that is identical to that for the normal configuration of the component under evaluation.

The testing of the mandatory controls shall include tests to demonstrate that the labels for information imported and/or exported to/from the component accurately represent the labels maintained by the NTCB partition for the component for use as the basis for its mandatory access

control decisions. The tests shall include each type of device, whether single-level or multilevel, supported by the component.

The NTCB must be FOUND RESISTANT to penetration. This applies to the NTCB as a whole, and to each NTCB partition in a component of this class.

+ Rationale

The phrase "no subject (without authorization to do so) is able to cause the TCB to enter a state such that it is unable to respond to communications initiated by other users" relates to the security services (Part II of this TNI) for the Denial of Service problem, and to correctness of the protocol implementations.

Testing is an important method available in this evaluation division to gain any assurance that the security mechanisms perform their intended function. A major purpose of testing is to demonstrate the system's response to inputs to the NTCB partition from untrusted (and possibly malicious) subjects.

In contrast to general purpose systems that allow for the dynamic creation of new programs and the introductions of new processes (and hence new subjects) with user specified security properties, many network components have no method for introducing new programs and/or processes during their normal operation. Therefore, the programs necessary for the testing must be introduced as special versions of the software rather than as the result of normal inputs by the test team. However, it must be insured that the NTCB partition used for such tests is identical to the one under evaluation.

Sensitivity labels serve a critical role in maintaining the security of the mandatory access controls in the network. Especially important to network security is the role of the labels for information communicated between components - explicit labels for multilevel devices and implicit labels for single-level devices. Therefore the testing for correct labels is highlighted.

The requirement for testing to demonstrate consistency between the NTCB implementation and the DTLs is a straightforward extension of the TCSEC requirement into the context of a network system.

3.3.3.2.2 Design Specification and Verification

+ Statement from DoD 5200.28-STD

A formal model of the security policy supported by the TCB shall be maintained over the life cycle of the ADP system that is proven and demonstrated to be consistent with its axioms. A descriptive top-level specification (DTLS) of the TCB shall be maintained that completely and accurately describes the TCB in terms of exceptions, error messages, and effects. It shall be shown to be an accurate description of the TCB interface. A CONVINCING ARGUMENT SHALL BE GIVEN THAT THE DTLS IS CONSISTENT WITH THE MODEL.

+ Interpretation

The overall network security policy expressed in this model will provide the basis for the mandatory access control policy exercised by the NTCB over subjects and storage objects in the entire network. The policy will also be the basis for the discretionary access control policy exercised by the NTCB to control access of named users to named objects. Data integrity requirements addressing the effects of unauthorized MSM need not be included in this model. The overall network policy must be decomposed into policy elements that are allocated to appropriate components and used as the basis for the security policy model for those components.

The level of abstraction of the model, and the set of subjects and objects that are explicitly represented in the model, will be affected by the NTCB partitioning. Subjects and objects must be represented explicitly in the model for the partition if there is some network component whose NTCB partition exercises access control over them. The model shall be structured so that the axioms and entities applicable to individual network components are manifest. Global network policy elements that are allocated to components shall be represented by the model for that component.

The requirements for a network DTLS are given in the Design Documentation section.

+ Rationale

The treatment of the model depends to a great extent on the degree of integration of the communications service into a distributed system. In a closely coupled distributed system, one might use a model that closely resembles one appropriate for a stand-alone computer system.

In ALL cases, the model of each partition will be expected to show the role of the NTCB partition in each kind of component. It will most likely clarify the model, although not part of the model, to show access restrictions

implied by the system design; for example, subjects representing protocol entities might have access only to objects containing data units at the same layer of protocol. The allocation of subjects and objects to different protocol layers is a protocol design choice which need not be reflected in the security policy model.

3.3.3.2.3 Configuration Management

+ Statement from DoD 5200.28-STD

During development and maintenance of the TCB, a configuration management system shall be in place that maintains control of changes to the descriptive top-level specification, other design data, implementation documentation, source code, the running version of the object code, and test fixtures and documentation. The configuration management system shall assure a consistent mapping among all documentation and code associated with the current version of the TCB. Tools shall be provided for generation of a new version of the TCB from source code. Also available shall be tools for comparing a newly generated version with the previous TCB version in order to ascertain that only the intended changes have been made in the code that will actually be used as the new version of the TCB.

+ Interpretation

The requirement applies as written, with the following extensions:

1. A configuration management system must be in place for each NTCB partition.
2. A configuration management plan must exist for the entire system. If the configuration management system is made up of the conglomeration of the configuration management systems of the various NTCB partitions, then the configuration management plan must address the issue of how configuration control is applied to the system as a whole.

+ Rationale

Each NTCB partition must have a configuration management system in place, or else there will be no way for the NTCB as a whole to have an effective configuration management system. The other extensions are merely reflections of the way that networks operate in practice.

3.3.4 Documentation.

3.3.4.1 Security Features User's Guide

+ Statement from DoD 5200.28-STD

A single summary, chapter, or manual in user documentation shall describe the protection mechanisms provided by the TCB, interpretations on their use, and how they interact with one another.

+ Interpretation

This user documentation describes user visible protection mechanisms at the global (network system) level and at the user interface of each component, and the interaction among these.

+ Rationale

The interpretation is an extension of the requirement into the context of a network system as defined for these network criteria. Documentation of protection mechanisms provided by individual components is required by the criteria for trusted computer systems that are applied as appropriate for the individual components.

3.3.4.2 Trusted Facility Manual

+ Statement from DoD 5200.28-STD

A manual addressed to the ADP system administrator shall present cautions about functions and privileges that should be controlled when running a secure facility. The procedures for examining and maintaining the audit files as well as the detailed audit record structure for each type of audit event shall be given. The manual shall describe the operator and administrator functions related to security, to include changing the security characteristics of a user. It shall provide interpretations on the consistent and effective use of the protection features of the system, how they interact, how to securely generate a new TCB, and facility procedures, warnings, and privileges that need to be controlled in order to operate the facility in a secure manner. The TCB modules that contain the reference validation mechanism shall be identified. The procedures for secure generation of a new TCB from source after modification of any modules in the TCB shall be described. IT SHALL INCLUDE THE PROCEDURES TO ENSURE THAT THE SYSTEM IS INITIALLY STARTED IN A SECURE

MANNER. PROCEDURES SHALL ALSO BE INCLUDED TO RESUME SECURE SYSTEM OPERATION AFTER ANY LAPSE IN SYSTEM OPERATION.

+ Interpretation

This manual shall contain specifications and procedures to assist the system administrator(s) maintain cognizance of the network configuration. These specifications and procedures shall address the following:

1. The hardware configuration of the network itself;
2. The implications of attaching new components to the network;
3. The case where certain components may periodically leave the network (e.g., by crashing, or by being disconnected) and then rejoin;
4. Network configuration aspects that can impact the security of the network system; (For example, the manual should describe for the network system administrator the interconnections among components that are consistent with the overall network system architecture.)
5. Loading or modifying NTCB software or firmware (e.g., down-line loading).
6. Incremental updates; that is, it must explicitly indicate which components of the network may change without others also changing.

The physical and administrative environmental controls shall be specified. Any assumptions about security of a given network should be clearly stated (e.g., the fact that all communications links must be physically protected to a certain level).

The components of the network that form the NTCB must be identified. Furthermore, the modules within an NTCB partition that contain the reference validation mechanism (if any) within that partition must be identified.

The procedures for the secure generation of a new version (or copy) of each NTCB partition from source must be described. The procedures and requirements for the secure generation of the NTCB necessitated by changes in the network configuration shall be described.

PROCEDURES FOR STARTING EACH NTCB PARTITION IN A SECURE STATE SHALL BE SPECIFIED. PROCEDURES MUST ALSO BE INCLUDED TO RESUME SECURE OPERATION OF EACH NTCB PARTITION AND/OR THE NTCB AFTER ANY LAPSE IN SYSTEM OR SUBSYSTEM OPERATION.

+ Rationale

There may be multiple system administrators with diverse responsibilities. The technical security measures described by these criteria must be used in conjunction with other forms of security in order to achieve security of the network. Additional forms include administrative security, physical security, emanations security, etc.

Extension of this criterion to cover configuration aspects of the network is needed because, for example, proper interconnection of components is typically essential to achieve a correct realization of the network architecture.

As mentioned in the section on Label Integrity, cryptography is one common mechanism employed to protect communication circuits. Encryption transforms the representation of information so that it is unintelligible to unauthorized subjects. Reflecting this transformation, the sensitivity of the ciphertext is generally lower than the cleartext. If encryption methodologies are employed, they shall be approved by the National Security Agency (NSA).

The encryption algorithm and its implementation are outside the scope of these interpretations. This algorithm and implementation may be implemented in a separate device or may be a function of a subject in a component not dedicated to encryption. Without prejudice, either implementation packaging is referred to as an encryption mechanism herein.

The requirements for descriptions of NTCB generation and identification of modules and components that form the NTCB are straightforward extensions of the TCSEC requirements into the network context. In those cases where the vendor does not provide source code, an acceptable procedure shall be to request the vendor to perform the secure generation.

GIVEN THE NATURE OF NETWORK SYSTEMS (E.G., VARIOUS COMPONENTS TEND TO BE DOWN AT DIFFERENT TIMES, AND THE NETWORK SYSTEM MUST CONTINUE OPERATION WITHOUT THAT COMPONENT), IT IS IMPERATIVE TO KNOW BOTH HOW TO SECURELY START UP AN NTCB PARTITION, AND HOW TO RESUME OPERATION SECURELY. IT IS ALSO NECESSARY TO KNOW HOW TO RESUME SECURE OPERATION OF THE NTCB AFTER ANY PARTITION HAS BEEN DOWN.

3.3.4.3 Test Documentation

+ Statement from DoD 5200.28-STD

The system developer shall provide to the evaluators a document that describes the test plan, test procedures that show how the security mechanisms were tested, and results of the security mechanisms' functional testing. It shall include results of testing the effectiveness of the methods used to reduce covert channel bandwidths.

+ Interpretation

The "system developer" is interpreted as "the network system sponsor". The description of the test plan should establish the context in which the testing was or should be conducted. The description should identify any additional test components that are not part of the system being evaluated. This includes a description of the test-relevant functions of such test components and a description of the interfacing of those test components to the system being evaluated. The description of the test plan should also demonstrate that the tests adequately cover the network security policy. The tests should include the features described in the System Architecture and the System Integrity sections. The tests should also include network configuration and sizing.

+ Rationale

The entity being evaluated may be a networking subsystem (see Appendix A) to which other components must be added to make a complete network system. In that case, this interpretation is extended to include contextual definition because, at evaluation time, it is not possible to validate the test plans without the description of the context for testing the networking subsystem.

The bandwidths of covert channels are used to determine the suitability of a network system for a given environment. The effectiveness of the methods used to reduce these bandwidths must therefore be accurately determined.

3.3.4.4 Design Documentation

+ Statement from DoD 5200.28-STD

Documentation shall be available that provides a description of the manufacturer's philosophy of protection and an explanation of how this philosophy is translated into the TCB. The interfaces between the TCB modules shall be described. A formal description of the security policy model enforced by the TCB shall be available and an explanation provided to show that it is sufficient to enforce the security policy. The specific TCB protection mechanisms shall be identified and an explanation given to show that they satisfy the model. The descriptive top-level specification (DTLS) shall

be shown to be an accurate description of the TCB interface. Documentation shall describe how the TCB implements the reference monitor concept and give an explanation why it is tamper resistant, cannot be bypassed, and is correctly implemented. THE TCB IMPLEMENTATION (I.E., IN HARDWARE, FIRMWARE, AND SOFTWARE) SHALL BE INFORMALLY SHOWN TO BE CONSISTENT WITH THE DTLs. THE ELEMENTS OF THE DTLs SHALL BE SHOWN, USING INFORMAL TECHNIQUES, TO CORRESPOND TO THE ELEMENTS OF THE TCB. Documentation shall describe how the TCB is structured to facilitate testing and to enforce least privilege. This documentation shall also present the results of the covert channel analysis and the tradeoffs involved in restricting the channels. All auditable events that may be used in the exploitation of known covert storage channels shall be identified. The bandwidths of known covert storage channels, the use of which is not detectable by the auditing mechanisms, shall be provided. (See the Covert Channel Guideline section.)

+ Interpretation

Explanation of how the sponsor's philosophy of protection is translated into the NTCB shall include a description of how the NTCB is partitioned. The security policy also shall be stated. The description of the interfaces between the NTCB modules shall include the interface(s) between NTCB partitions and modules within the partitions if the modules exist. The sponsor shall describe the security architecture and design, including the allocation of security requirements among components.

The documentation includes both a system description and a set of component DTLs's. The system description addresses the network security architecture and design by specifying the types of components in the network, which ones are trusted, and in what way they must cooperate to support network security objectives. A component DTLs shall be provided for each trusted network component, i.e., each component containing an NTCB partition. Each component DTLs shall describe the interface to the NTCB partition of its component. BOTH THE SYSTEM DESCRIPTION AND EACH COMPONENT DTLs SHALL BE SHOWN CONSISTENT WITH THOSE ASSERTIONS IN THE MODEL THAT APPLY TO IT. Appendix A addresses component evaluation issues.

TO SHOW THE CORRESPONDENCE BETWEEN THE DTLs AND THE NTCB IMPLEMENTATION, IT SUFFICES TO SHOW CORRESPONDENCE BETWEEN EACH COMPONENT DTLs AND THE NTCB PARTITION IN THAT COMPONENT.

As stated in the introduction to Division B, the sponsor must demonstrate that the NTCB employs the reference monitor concept. The security policy model must be a model for a reference monitor.

The security policy model for each partition implementing a reference monitor shall fully represent the access control policy supported by the partition, including the discretionary and mandatory security policy for secrecy and/or integrity. For the mandatory policy the single dominance relation for sensitivity labels, including secrecy and/or integrity components, shall be precisely defined.

+ Rationale

The interpretation is a straightforward extension of the requirement into the context of a network system as defined for this network interpretation. Other documentation, such as description of components and description of operating environment(s) in which the networking subsystem or network system is designed to function, is required elsewhere, e.g., in the Trusted Facility Manual.

In order to be evaluated, a network must possess a coherent Network Security Architecture and Design. (Interconnection of components that do not adhere to such a single coherent Network Security Architecture is addressed in the Interconnection of Accredited AIS, Appendix C.) The Network Security Architecture must address the security-relevant policies, objectives, and protocols. The Network Security Design specifies the interfaces and services that must be incorporated into the network so that it can be evaluated as a trusted entity. There may be multiple designs that conform to the same architecture but are more or less incompatible and non-interoperable (except through the Interconnection Rules). Security related mechanisms requiring cooperation among components are specified in the design in terms of their visible interfaces; mechanisms having no visible interfaces are not specified in this document but are left as implementation decisions.

The Network Security Architecture and Design must be available from the network sponsor before evaluation of the network, or any component, can be undertaken. The Network Security Architecture and Design must be sufficiently complete, unambiguous, and free from obvious flaws to permit the construction or assembly of a trusted network based on the structure it specifies.

When a component is being designed or presented for evaluation, or when a network assembled from components is assembled or presented for evaluation, there must be a priori evidence that the Network security Architecture and Design are satisfied. That is, the components can be assembled into a network that conforms in every way with the Network Security Architecture and Design to produce a physical realization that is trusted to the extent that its evaluation indicates.

In order for a trusted network to be constructed from components that can be built independently, the Network Security Architecture and Design must completely and unambiguously define the security functionality of components as well as the interfaces between or among components. The Network Security Architecture and Design must be evaluated to determine that a network constructed to its specifications will in fact be trusted, that is, it will be evaluable under these interpretations.

The term "model" is used in several different ways in a network context, e.g., a "protocol reference model," a "formal network model," etc. Only the "security policy model" is addressed by this requirement and is specifically intended to model the interface, viz., "security perimeter," of the reference monitor and must meet all the requirements defined in the TCSEC. It must be shown that all parts of the TCB are a valid interpretation of the security policy model, i.e., that there is no change to the secure state except as represented by the model.

4.0 DIVISION A: VERIFIED PROTECTION

This division is characterized by the use of formal security methods to assure that the mandatory and discretionary security controls employed in the network system can effectively protect classified or other sensitive information stored or processed by the system. Extensive documentation is required to demonstrate that the NTCB meets the security requirements in all aspects of design, development and implementation.

4.1 CLASS (A1): VERIFIED DESIGN

SYSTEMS IN CLASS (A1) ARE FUNCTIONALLY EQUIVALENT TO THOSE IN CLASS (B3) IN THAT NO ADDITIONAL ARCHITECTURAL FEATURES OR POLICY REQUIREMENTS ARE ADDED. THE DISTINGUISHING FEATURE OF SYSTEMS IN THIS CLASS IS THE ANALYSIS DERIVED FROM FORMAL DESIGN SPECIFICATION AND VERIFICATION TECHNIQUES AND THE RESULTING HIGH DEGREE OF ASSURANCE THAT THE NTCB IS CORRECTLY IMPLEMENTED. THIS ASSURANCE IS DEVELOPMENTAL IN NATURE, STARTING WITH A FORMAL MODEL OF THE SECURITY POLICY AND A FORMAL TOP-LEVEL SPECIFICATION (FTLS) OF THE DESIGN. INDEPENDENT OF THE PARTICULAR SPECIFICATION LANGUAGE OR VERIFICATION SYSTEM USED, THERE ARE

FIVE IMPORTANT CRITERIA FOR CLASS (A1) DESIGN VERIFICATION:

- + A FORMAL MODEL OF THE SECURITY POLICY MUST BE CLEARLY IDENTIFIED AND DOCUMENTED, INCLUDING A MATHEMATICAL PROOF THAT THE MODEL IS CONSISTENT WITH ITS AXIOMS AND IS SUFFICIENT TO SUPPORT THE SECURITY POLICY.
- + AN FTLS MUST BE PRODUCED THAT INCLUDES ABSTRACT DEFINITIONS OF THE FUNCTIONS THE NTCB PERFORMS AND OF THE HARDWARE AND/OR FIRMWARE MECHANISMS THAT ARE USED TO SUPPORT SEPARATE EXECUTION DOMAINS.
- + THE FTLS OF THE NTCB MUST BE SHOWN TO BE CONSISTENT WITH THE MODEL BY FORMAL TECHNIQUES WHERE POSSIBLE (I.E., WHERE VERIFICATION TOOLS EXIST) AND INFORMAL ONES OTHERWISE.
- + THE NTCB IMPLEMENTATION (I.E., IN HARDWARE, FIRMWARE, AND SOFTWARE) MUST BE INFORMALLY SHOWN TO BE CONSISTENT WITH THE FTLS. THE ELEMENTS OF THE FTLS MUST BE SHOWN, USING INFORMAL TECHNIQUES, TO CORRESPOND TO THE ELEMENTS OF THE NTCB. THE FTLS MUST EXPRESS THE UNIFIED PROTECTION MECHANISM REQUIRED TO SATISFY THE SECURITY POLICY, AND IT IS THE ELEMENTS OF THIS PROTECTION MECHANISM THAT ARE MAPPED TO THE ELEMENTS OF THE NTCB.
- + FORMAL ANALYSIS TECHNIQUES MUST BE USED TO IDENTIFY AND ANALYZE COVERT CHANNELS. INFORMAL TECHNIQUES MAY BE USED TO IDENTIFY COVERT TIMING CHANNELS. THE CONTINUED EXISTENCE OF IDENTIFIED COVERT CHANNELS IN THE SYSTEM MUST BE JUSTIFIED.

IN KEEPING WITH THE EXTENSIVE DESIGN AND DEVELOPMENT ANALYSIS OF THE NTCB REQUIRED OF SYSTEMS IN CLASS (A1), MORE STRINGENT CONFIGURATION MANAGEMENT IS REQUIRED AND PROCEDURES ARE ESTABLISHED FOR SECURELY DISTRIBUTING THE SYSTEM TO SITES. A SYSTEM SECURITY ADMINISTRATOR IS SUPPORTED.

THE FOLLOWING ARE MINIMAL REQUIREMENTS FOR SYSTEM ASSIGNED A CLASS (A1) RATING:

4.1.1 Security Policy

- + Statement from DoD 5200.28-STD

Implied from the Introduction to the TCSEC.

+ Interpretation

The network sponsor shall describe the overall network security policy enforced by the NTCB. At a minimum, this policy shall include the discretionary and mandatory requirements applicable to this class. The policy may require data secrecy, or data integrity, or both. The policy is an access control policy having two primary components: mandatory and discretionary. The policy shall include a discretionary policy for protecting the information being processed based on the authorizations of individuals, users, or groups of users. This access control policy statement shall describe the requirements on the network to prevent or detect "reading or destroying" sensitive information by unauthorized users or errors. The mandatory policy must define the set of distinct sensitivity levels that it supports. For the Class B1 or above the mandatory policy shall be based on the labels associated with the information that reflects its sensitivity with respect to secrecy and/or integrity, where applicable, and labels associated with users to reflect their authorization to access such information. Unauthorized users include both those that are not authorized to use the network at all (e.g., a user attempting to use a passive or active wire tap) or a legitimate user of the network who is not authorized to access a specific piece of information being protected.

Note that "users" does not include "operators," "system programmers," "technical control officers," "system security officers," and other system support personnel. They are distinct from users and are subject to the Trusted Facility Manual and the System Architecture requirements. Such individuals may change the system parameters of the network system, for example, by defining membership of a group. These individuals may also have the separate role of users.

SECRECY POLICY: The network sponsor shall define the form of the discretionary and mandatory secrecy policy that is enforced in the network to prevent unauthorized users from reading the sensitive information entrusted to the network.

DATA INTEGRITY POLICY: The network sponsor shall define the discretionary and mandatory integrity policy to prevent unauthorized users from modifying, viz., writing, sensitive information. The definition of data integrity presented by the network sponsor refers to the requirement that the information has not been subjected to unauthorized modification in the network. The mandatory integrity pol-

icy enforced by the NTCB cannot, in general, prevent modification while information is being transmitted between components. However, an integrity sensitivity label may reflect the confidence that the information has not been subjected to transmission errors because of the protection afforded during transmission. This requirement is distinct from the requirement for label integrity.

+ Rationale

The word "sponsor" is used in place of alternatives (such as "vendor," "architect," "manufacturer," and "developer") because the alternatives indicate people who may not be available, involved, or relevant at the time that a network system is proposed for evaluation.

A trusted network is able to control both the reading and writing of shared sensitive information. Control of writing is used to protect against destruction of information. A network normally is expected to have policy requirements to protect both the secrecy and integrity of the information entrusted to it. In a network the integrity is frequently as important or more important than the secrecy requirements. Therefore the secrecy and/or integrity policy to be enforced by the network must be stated for each network regardless of its evaluation class. The assurance that the policy is faithfully enforced is reflected in the evaluation class of the network.

This control over modification is typically used to protect information so that it may be relied upon and to control the potential harm that would result if the information were corrupted. The overall network policy requirements for integrity includes the protection for data both while being processed in a component and while being transmitted in the network. The access control policy enforced by the NTCB relates to the access of subjects to objects within each component. Communications integrity addressed within Part II relates to information while being transmitted.

The mandatory integrity policy (at class B1 and above) in some architectures may be useful in supporting the linkage between the connection oriented abstraction introduced in the Introduction and the individual components of the network. For example, in a key distribution center for end-to-end encryption, a distinct integrity category may be assigned to isolate the key generation code and data from possible modification by other supporting processes in the same component, such as operator interfaces and audit.

The mandatory integrity policy for some architecture may define an integrity sensitivity label that reflects the

specific requirements for ensuring that information has not been subject to random errors in excess of a stated limit nor to unauthorized message stream modification (MSM) -. The specific metric associated with an integrity sensitivity label will generally reflect the intended applications of the network.

4.1.1.1 Discretionary Access Control

+ Statement from DoD 5200.28-STD

The TCB shall define and control access between named users and named objects (e.g., files and programs) in the ADP system. The enforcement mechanism (e.g., access control lists) shall allow users to specify and control sharing of those objects and shall provide controls to limit propagation of access rights. The discretionary access control mechanism shall, either by explicit user action or by default, provide that objects are protected from unauthorized access. These access controls shall be capable of specifying, for each named object, a list of named individuals and a list of groups of named individuals with their respective modes of access to that object. Furthermore, for each such named object, it shall be possible to specify a list of named individuals and a list of groups of named individuals for which no access to the object is given. Access permission to an object by users not already possessing access

- See Voydock, Victor L. and Stephen T. Kent, "Security Mechanisms in High-Level Network Protocols," Computing Surveys, Vol. 15, No. 2, June 1983, pp 135-171.

permission shall only be assigned by authorized users.

+ Interpretation

The discretionary access control (DAC) mechanism(s) may be distributed over the partitioned NTCB in various ways. Some part, all, or none of the DAC may be implemented in a given component of the network system. In particular, components that support only internal subjects (i.e., that have no subjects acting as direct surrogates for users), such as a public network packet switch, might not implement the DAC mechanism(s) directly (e.g., they are unlikely to contain access control lists).

Identification of users by groups may be achieved in various ways in the networking environment. For example, the network identifiers (e.g., internet addresses) for various components (e.g., hosts, gateways) can be used as iden-

tifiers of groups of individual users (e.g., "all users at Host A," "all users of network Q") so long as the individuals involved in the group are implied by the group identifier. For example, Host A might employ a particular group-id, for which it maintains a list of explicit users in that group, in its network exchange with Host B, which accepts the group-id under the conditions of this interpretation.

For networks, individual hosts will impose need-to-know controls over their users on the basis of named individuals - much like (in fact, probably the same) controls used when there is no network connection.

When group identifiers are acceptable for access control, the identifier of some other host may be employed, to eliminate the maintenance that would be required if individual identification of remote users was employed. In class C2 and higher, however, it must be possible from that audit record to identify (immediately or at some later time) exactly the individuals represented by a group identifier at the time of the use of that identifier. There is allowed to be an uncertainty because of elapsed time between changes in the group membership and the enforcement in the access control mechanisms.

The DAC mechanism of a NTCB partition may be implemented at the interface of the reference monitor or may be distributed in subjects that are part of the NTCB in the same or different component. The reference monitor manages all the physical resources of the system and from them creates the abstraction of subjects and objects that it controls. Some of these subjects and objects may be used to implement a part of the NTCB. When the DAC mechanism is distributed in such NTCB subjects (i.e., when outside the reference monitor), the assurance requirements (see the Assurance section) for the design and implementation of the DAC shall be those of class C2 for all networks of class C2 or above.

When integrity is included as part of the network discretionary security policy, the above interpretations shall be specifically applied to the controls over modification, viz, the write mode of access, within each component based on identified users or groups of users.

+ Rationale

In this class, the supporting elements of the overall DAC mechanism are required to isolate information (objects) that supports DAC so that it is subject to auditing requirements (see the System Architecture section). The use of network identifiers to identify groups of individual users could be implemented, for example, as an X.25 community of interest in the network protocol layer (layer 3). In all

other respects, the supporting elements of the overall DAC mechanism are treated exactly as untrusted subjects are treated with respect to DAC in an ADP system, with the same result as noted in the interpretation.

A typical situation for DAC is that a surrogate process for a remote user will be created in some host for access to objects under the control of the NTCB partition within that host. The interpretation requires that a user identifier be assigned and maintained for each such process by the NTCB, so that access by a surrogate process is subject to essentially the same discretionary controls as access by a process acting on behalf of a local user would be. However, within this interpretation a range of possible interpretations of the assigned user identification is permitted.

The most obvious situation would exist if a global database of network users were to be made permanently available on demand to every host, (i.e., a name server existed) so that all user identifications were globally meaningful.

It is also acceptable, however, for some NTCB partitions to maintain a database of locally-registered users for its own use. In such a case, one could choose to inhibit the creation of surrogate processes for locally unregistered users, or (if permitted by the local policy) alternatively, to permit the creation of surrogate processes with preselected user and group identifiers which, in effect, identify the process as executing on behalf of a member of a group of users on a particular remote host. The intent of the words concerning audit in the interpretation is to provide a minimally acceptable degree of auditability for cases such as the last described. What is required is that there be a capability, using the audit facilities provided by the network NTCB partitions involved, to determine who was logged in at the actual host of the group of remote users at the time the surrogate processing occurred.

Associating the proper user id with a surrogate process is the job of identification and authentication. This means that DAC is applied locally, with respect to the user id of the surrogate process. The transmission of the data back across the network to the user's host, and the creation of a copy of the data there, is not the business of DAC.

Components that support only internal subjects impact the implementation of the DAC by providing services by which information (e.g., a user-id) is made available to a component that makes a DAC decision. An example of the latter would be the case that a user at Host A attempts to access a file at Host B. The DAC decision might be (and usually would be) made at Host B on the basis of a user-id transmitted from Host A to Host B.

Unique user identification may be achieved by a variety of mechanisms, including (a) a requirement for unique identification and authentication on the host where access takes place; (b) recognition of fully qualified network addresses authenticated by another host and forwarded to the host where access takes place; or (c) administrative support of a network-wide unique personnel identifier that could be authenticated and forwarded by another host as in (b) above, or could be authenticated and forwarded by a dedicated network identification and authentication server. The protocols which implement (b) or (c) are subject to the System Architecture requirements.

Network support for DAC might be handled in other ways than that described as "typical" above. In particular, some form of centralized access control is often proposed. An access control center may make all decisions for DAC, or it may share the burden with the hosts by controlling host-to-host connections, and leaving the hosts to decide on access to their objects by users at a limited set of remote hosts. In this case the access control center provides the linkage between the connection oriented abstraction (as discussed in the Introduction) and the overall network security policy for DAC. In all cases the enforcement of the decision must be provided by the host where the object resides.

There are two forms of distribution for the DAC mechanism: implementing portions of the DAC in separate components, and supporting the DAC in subjects contained within the NTCB partition in a component. Since "the ADP system" is understood to be "the computer network" as a whole, each network component is responsible for enforcing security in the mechanisms allocated to it to ensure secure implementation of the network security policy. For traditional host systems it is frequently easy to also enforce the DAC along with the MAC within the reference monitor, per se, although a few approaches, such as virtual machine monitors, support DAC outside this interface.

In contrast to the universally rigid structure of mandatory policies (see the Mandatory Access Control section), DAC policies tend to be very network and system specific, with features that reflect the natural use of the system. For networks it is common that individual hosts will impose controls over their local users on the basis of named individuals--much like the controls used when there is no network connection. However, it is difficult to manage in a centralized manner all the individuals using a large network. Therefore, users on other hosts are commonly grouped together so that the controls required by the network DAC policy are actually based on the identity of the hosts or other components. A gateway is an example of such a component.

The assurance requirements are at the very heart of the concept of a trusted system. It is the assurance that determines if a system or network is appropriate for a given environment, as reflected, for example, in the Environments Guideline-. In the case of monolithic systems that have DAC integral to the reference monitor, the assurance requirements for DAC are inseparable from those of the rest of the reference monitor. For networks there is typically a much clearer distinction due to distributed DAC. The rationale for making the distinction in this network interpretation is that if major trusted network components can be made significantly easier to design and implement without reducing the ability to meet security policy, then trusted networks will be more easily available.

4.1.1.2 Object Reuse

+ Statement from DoD 5200.28-STD

All authorizations to the information contained within a storage object shall be revoked prior to initial assignment, allocation or reallocation to a subject from the TCB's pool of unused storage objects. No information, including encrypted representations of information, produced by a prior subject's actions is to be available to any subject that obtains access to an object that has been released back to the system.

+ Interpretation

The NTCB shall ensure that any storage objects that it controls (e.g., message buffers under the control of a NTCB partition in a component) contain no information for which a subject in that component is not authorized before granting access. This requirement must be enforced by each of the NTCB partitions.

- Guidance for Applying the Department of Defense
Trusted Computer System Evaluation Criteria in Specific
Environments, CSC-STD-003-85.

+ Rationale

In a network system, storage objects of interest are things that the NTCB directly controls, such as message buffers in components. Each component of the network system

must enforce the object reuse requirement with respect to the storage objects of interest as determined by the network security policy. For example, the DAC requirement in this division leads to the requirement here that message buffers be under the control of the NTCB partition. A buffer assigned to an internal subject may be reused at the discretion of that subject which is responsible for preserving the integrity of message streams. Such controlled objects may be implemented in physical resources, such as buffers, disk sectors, tape space, and main memory, in components such as network switches.

4.1.1.3 Labels

+ Statement from DoD 5200.28-STD

Sensitivity labels associated with each ADP system resource (e.g., subject, storage object, ROM) that is directly or indirectly accessible by subjects external to the TCB shall be maintained by the TCB. These labels shall be used as the basis for mandatory access control decisions. In order to import non-labeled data, the TCB shall request and receive from an authorized user the sensitivity level of the data, and all such actions shall be auditable by the TCB.

+ Interpretation

Non-labeled data imported under the control of the NTCB partition will be assigned a label constrained by the device labels of the single-level device used to import it. Labels may include secrecy and integrity- components in accordance with the overall network security policy described by the network sponsor. Whenever the term "label" is used throughout this interpretation, it is understood to include both components as applicable. Similarly, the terms "single-level" and "multilevel" are understood to be based on both the secrecy and integrity components of the policy. The mandatory integrity policy will typically have requirements, such as the probability of undetected message stream modification, that will be reflected in the label for the data so protected. For example, when data is imported its integrity label may be assigned based on mechanisms, such as cryptography, used to provide the assurance required by the policy. The NTCB shall assure that such mechanism are protected from tampering and are always invoked when they are

- See, for example, Biba, K.J., "Integrity Consideration for Secure Computer Systems," ESD-TR-76-372, MTR-3153, The MITRE Corporation, Bedford, MA, April 1977.

the basis for a label.

If the security policy includes an integrity policy, all activities that result in message-stream modification during transmission are regarded as unauthorized accesses in violation of the integrity policy. The NTCB shall have an automated capability for testing, detecting, and reporting those errors/corruptions that exceed specified network integrity policy requirements. Message-stream modification (MSM) countermeasures shall be identified. A technology of adequate strength shall be selected to resist MSM. If encryption methodologies are employed, they shall be approved by the National Security Agency.

All objects must be labeled within each component of the network that is trusted to maintain separation of multiple levels of information. The label associated with any objects associated with single-level components will be identical to the level of that component. Objects used to store network control information, and other network structures, such as routing tables, must be labeled to prevent unauthorized access and/or modification.

+ Rationale

The interpretation is an extension of the requirement into the context of a network system and partitioned NTCB as defined for these network interpretations. A single-level device may be regarded either as a subject or an object. A multilevel device is regarded as a trusted subject in which the security range of the subject is the minimum-maximum range of the data expected to be transmitted over the device.

The sensitivity labels for either secrecy or integrity or both may reflect non-hierarchical categories or hierarchical classification or both.

For a network it is necessary that this requirement be applied to all network system resources at the (B2) level and above.

The NTCB is responsible for implementing the network integrity policy, when one exists. The NTCB must enforce that policy by ensuring that information is accurately transmitted from source to destination (regardless of the number of intervening connecting points). The NTCB must be able to counter equipment failure, environmental disruptions, and actions by persons and processes not authorized to alter the data. Protocols that perform code or format conversion shall preserve the integrity of data and control information.

The probability of an undetected transmission error may

be specified as part of the network security policy so that the acceptability of the network for its intended application may be determined. The specific metrics (e.g., probability of undetected modification) satisfied by the data can be reflected in the integrity sensitivity label associated with the data while it is processed within a component. It is recognized that different applications and operational environments (e.g., crisis as compared to logistic) will have different integrity requirements.

The network shall also have an automated capability of testing for, detecting, and reporting errors that exceed a threshold consistent with the operational mode requirements. The effectiveness of integrity countermeasures must be established with the same rigor as the other security-relevant properties such as secrecy.

Cryptography is often utilized as a basis to provide data integrity assurance. Mechanisms, such as Manipulation Detection Codes (MDC)-, may be used. The adequacy of the encryption or MDC algorithm, the correctness of the protocol logic, and the adequacy of implementation must be established in MSM countermeasures design.

4.1.1.3.1 Label Integrity

+ Statement from DoD 5200.28-STD

Sensitivity labels shall accurately represent sensitivity levels of the specific subjects or objects with which they are associated. When exported by the TCB, sensitivity labels shall accurately and unambiguously represent the internal labels and shall be associated with the information being exported.

+ Interpretation

The phrase "exported by the TCB" is understood to include transmission of information from an object in one component to an object in another component. Information transferred between NTCB partitions is addressed in the System Integrity Section. The form of internal and external (exported) sensitivity labels may differ, but the meaning shall be the same. The NTCB shall, in addition, ensure that correct association of sensitivity labels with the information being transported across the network is preserved.

As mentioned in the Trusted Facility Manual Section, encryption transforms the representation of information so that it is unintelligible to unauthorized subjects. Reflecting this transformation, the sensitivity level of the ciphertext is generally lower than the cleartext. It fol-

lows that cleartext and ciphertext are contained in

- See Jueneman, R. R., "Electronic Document Authentication," IEEE Network Magazine, April 1987, pp 17-23.

different objects, each possessing its own label. The label of the cleartext must be preserved and associated with the ciphertext so that it can be restored when the cleartext is subsequently obtained by decrypting the ciphertext. If the cleartext is associated with a single-level device, the label of that cleartext may be implicit. The label may also be implicit in the key.

When information is exported to an environment where it is subject to deliberate or accidental modification, the TCB shall support the means, such as cryptographic checksums, to assure the accuracy of the labels. When there is a mandatory integrity policy, the policy will define the meaning of integrity labels.

+ Rationale

Encryption algorithms and their implementation are outside the scope of these interpretations. Such algorithms may be implemented in a separate device or may be incorporated in a subject of a larger component. Without prejudice, either implementation packaging is referred to as an encryption mechanism herein. If encryption methodologies are employed in this regard, they shall be approved by the National Security Agency (NSA). The encryption process is part of the Network Trusted Computer Base partition in the components in which it is implemented.

The encryption mechanism is not necessarily a multilevel device or multilevel subject, as these terms are used in these criteria. The process of encryption is multilevel by definition. The cleartext and ciphertext interfaces carry information of different sensitivity. An encryption mechanism does not process data in the sense of performing logical or arithmetic operations on that data with the intent of producing new data. The cleartext and ciphertext interfaces on the encryption mechanism must be separately identified as being single-level or multilevel. If the interface is single-level, then the sensitivity of the data is established by a trusted individual and implicitly associated with the interface; the Exportation to Single-Level Devices criterion applies.

If the interface is multilevel, then the data must be labeled; the Exportation to Multilevel Devices criterion applies. The network architect is free to select an acceptable mechanism for associating a label with an object. With

reference to encrypted objects, the following examples are possible:

1. Include a label field in the protocol definition of the object.
2. Implicitly associate the label with the object through the encryption key. That is, the encryption key uniquely identifies a sensitivity level. A single or private key must be protected at the level of the data that it encrypts.

4.1.1.3.2 Exportation of Labeled Information

+ Statement from DoD 5200.28-STD

The TCB shall designate each communication channel and I/O device as either single-level or multilevel. Any change in this designation shall be done manually and shall be auditable by the TCB. The TCB shall maintain and be able to audit any change in the sensitivity level or levels associated with a communications channel or I/O device.

+ Interpretation

Each communication channel and network component shall be designated as either single-level or multilevel. Any change in this designation shall be done with the cognizance and approval of the administrator or security officer in charge of the affected components and the administrator or security officer in charge of the NTCB. This change shall be auditable by the network. The NTCB shall maintain and be able to audit any change in the device labels associated with a single-level communication channel or the range associated with a multilevel communication channel or component. The NTCB shall also be able to audit any change in the set of sensitivity levels associated with the information which can be transmitted over a multilevel communication channel or component.

+ Rationale

Communication channels and components in a network are analogous to communication channels and I/O devices in stand-alone systems. They must be designated as either multilevel (i.e., able to distinguish and maintain separation among information of various sensitivity levels) or single-level. As in the TCSEC, single-level devices may only be attached to single-level channels.

The level or set of levels of information that can be sent to a component or over a communication channel shall only change with the knowledge and approval of the security

officers (or system administrator, if there is no security officer) of the network, and of the affected components. This requirement ensures that no significant security-relevant changes are made without the approval of all affected parties.

4.1.1.3.2.1 Exportation to Multilevel Devices

+ Statement from DoD 5200.28-STD

When the TCB exports an object to a multilevel I/O device, the sensitivity label associated with that object shall also be exported and shall reside on the same physical medium as the exported information and shall be in the same form (i.e., machine-readable or human-readable form). When the TCB exports or imports an object over a multilevel communications channel, the protocol used on that channel shall provide for the unambiguous pairing between the sensitivity labels and the associated information that is sent or received.

+ Interpretation

The components, including hosts, of a network shall be interconnected over "multilevel communication channels," multiple single-level communication channels, or both, whenever the information is to be protected at more than a single sensitivity level. The protocol for associating the sensitivity label and the exported information shall provide the only information needed to correctly associate a sensitivity level with the exported information transferred over the multilevel channel between the NTCB partitions in individual components. This protocol definition must specify the representation and semantics of the sensitivity labels (i.e., the machine-readable label must uniquely represent the sensitivity level).

The "unambiguous" association of the sensitivity level with the communicated information shall meet the same level of accuracy as that required for any other label within the NTCB, as specified in the criterion for Label Integrity. This may be provided by protected and highly reliable direct physical layer connections, or by traditional cryptographic link protection in which any errors during transmission can be readily detected, or by use of a separate channel. The range of information imported or exported must be constrained by the associated device labels.

+ Rationale

This protocol must specify the representation and semantics of the sensitivity labels. See the Mandatory Access Control Policies section in Appendix B. The mul-

tilevel device interface to (untrusted) subjects may be implemented either by the interface of the reference monitor, per se, or by a multilevel subject (e.g., a "trusted subject" as defined in the Bell-LaPadula Model) that provides the labels based on the internal labels of the NTCB partition.

The current state of the art limits the support for mandatory policy that is practical for secure networks. Reference monitor support to ensure the control over all the operations of each subject in the network must be completely provided within the single NTCB partition on which that subject interfaces to the NTCB. This means that the entire portion of the "secure state" represented in the formal security policy model that may be changed by transitions invoked by this subject must be contained in the same component.

The secure state of an NTCB partition may be affected by events external to the component in which the NTCB partition resides (e.g., arrival of a message). The effect occurs asynchronously after being initiated by an event in another component or partition. For example, indeterminate delays may occur between the initiation of a message in one component, the arrival of the message in the NTCB partition in another component, and the corresponding change to the secure state of the second component. Since each component is executing concurrently, to do otherwise would require some sort of network-wide control to synchronize state transitions, such as a global network-wide clock for all processors; in general, such designs are not practical and probably not even desirable. Therefore, the interaction between NTCB partitions is restricted to just communications between pairs (at least logically) of devices-multilevel devices if the device(s) can send/receive data of more than a single level. For broadcast channels the pairs are the sender and intended receiver(s). However, if the broadcast channel carries multiple levels of information, additional mechanism (e.g., cryptochecksum maintained by the TCB) may be required to enforce separation and proper delivery.

A common representation for sensitivity labels is needed in the protocol used on that channel and understood by both the sender and receiver when two multilevel devices (in this case, in two different components) are interconnected. Each distinct sensitivity level of the overall network policy must be represented uniquely in these labels.

Within a monolithic TCB, the accuracy of the sensitivity labels is generally assured by simple techniques, e.g., very reliable connections over very short physical connections, such as on a single printed circuit board or over an internal bus. In many network environments there is a much higher probability of accidentally or maliciously

introduced errors, and these must be protected against.

4.1.1.3.2.2 Exportation to Single-Level Devices

+ Statement from DoD 5200.28-STD

Single-level I/O devices and single-level communication channels are not required to maintain the sensitivity labels of the information they process. However, the TCB shall include a mechanism by which the TCB and an authorized user reliably communicate to designate the single sensitivity level of information imported or exported via single-level communication channels or I/O devices.

+ Interpretation

Whenever one or both of two directly connected components is not trusted to maintain the separation of information of different sensitivity levels, or whenever the two directly connected components have only a single sensitivity level in common, the two components of the network shall communicate over a single-level channel. Single-level components and single-level communication channels are not required to maintain the sensitivity labels of the information they process. However, the NTCB shall include a reliable communication mechanism by which the NTCB and an authorized user (via a trusted path) or a subject within an NTCB partition can designate the single sensitivity level of information imported or exported via single-level communication channels or network components. The level of information communicated must equal the device level.

+ Rationale

Single-level communications channels and single-level components in networks are analogous to single level channels and I/O devices in stand-alone systems in that they are not trusted to maintain the separation of information of different sensitivity levels. The labels associated with data transmitted over those channels and by those components are therefore implicit; the NTCB associates labels with the data because of the channel or component, not because of an explicit part of the bit stream. Note that the sensitivity level of encrypted information is the level of the ciphertext rather than the original level(s) of the plaintext.

4.1.1.3.2.3 Labeling Human-Readable Output

+ Statement from DoD 5200.28-STD

The ADP system administrator shall be able to specify the printable label names associated with exported sensitivity

labels. The TCB shall mark the beginning and end of all human-readable, paged, hardcopy output (e.g., line printer output) with human-readable sensitivity labels that properly¹ represent the sensitivity of the output. The TCB shall, by default, mark the top and bottom of each page of human-readable, paged, hardcopy output (e.g., line printer output) with human-readable sensitivity labels that properly¹ represent the sensitivity of the page. The TCB shall, by default and in an appropriate manner, mark other forms of human readable output (e.g., maps, graphics) with human-readable sensitivity labels that properly¹ represent the sensitivity of the output. Any override of these markings defaults shall be auditable by the TCB.

¹ The hierarchical classification component in human-readable sensitivity labels shall be equal to the greatest hierarchical classification of any of the information in the output that the labels refer to; the non-hierarchical category component shall include all of the non-hierarchical categories of the information in the output the labels refer to, but no other non-hierarchical categories.

+ Interpretation

This criterion imposes no requirement to a component that produces no human-readable output. For those that do produce human-readable output, each sensitivity level that is defined to the network shall have a uniform meaning across all components. The network administrator, in conjunction with any affected component administrator, shall be able to specify the human-readable label that is associated with each defined sensitivity level.

+ Rationale

The interpretation is a straightforward extension of the requirement into the context of a network system and partitioned NTCB as defined for these network interpretations.

4.1.1.3.3 Subject Sensitivity Labels

+ Statement from DoD 5200.28-STD

The TCB shall immediately notify a terminal user of each change in the sensitivity level associated with that user during an interactive session. A terminal user shall be able to query the TCB as desired for a display of the subject's complete sensitivity label.

+ Interpretation

An NTCB partition shall immediately notify a terminal user attached to its component of each change in the sensitivity level associated with that user.

+ Rationale

The local NTCB partition must ensure that the user understands the sensitivity level of information sent to and from a terminal. When a user has a surrogate process in another component, adjustments to its level may occur to maintain communication with the user. These changes may occur asynchronously. Such adjustments are necessitated by mandatory access control as applied to the objects involved in the communication path.

4.1.1.3.4 Device Labels

+ Statement from DoD 5200.28-STD

The TCB shall support the assignment of minimum and maximum sensitivity levels to all attached physical devices. These sensitivity levels shall be used by the TCB to enforce constraints imposed by the physical environments in which the devices are located.

+ Interpretation

This requirement applies as written to each NTCB partition that is trusted to separate information based on sensitivity level. Each I/O device in a component, used for communication with other network components, is assigned a device range, consisting of a set of labels with a maximum and minimum. (A device range usually contains, but does not necessarily contain, all possible labels "between" the maximum and minimum, in the sense of dominating the minimum and being dominated by the maximum.)

The NTCB always provides an accurate label for information exported through devices. Information exported or imported using a single-level device is labelled implicitly by the sensitivity level of the device. Information exported from one multilevel device and imported at another must be labelled through an agreed-upon protocol, unless it is labelled implicitly by using a communication link that always carries a single level.

Information exported at a given sensitivity level can be sent only to an importing device whose device range contains that level or a higher level. If the importing device range does not contain the given level, the information is relabelled upon reception at a higher level within the importing device range. Relabelling should not occur otherwise.

+ Rationale

The purpose of device labels is to reflect and constrain the sensitivity levels of information authorized for the physical environment in which the devices are located.

The information transfer restrictions permit one-way communication (i.e., no acknowledgements) from one device to another whose ranges have no level in common, as long as each level in the sending device range is dominated by some level in the receiving device range. It is never permitted to send information at a given level to a device whose range does not contain a dominating level. (See Appendix C for similar interconnection rules for the interconnected AIS view.)

4.1.1.4 Mandatory Access Control

+ Statement from DoD 5200.28-STD

The TCB shall enforce a mandatory access control policy over all resources (i.e., subjects, storage objects, and I/O devices) that are directly or indirectly accessible by subjects external to the TCB. These subjects and objects shall be assigned sensitivity labels that are a combination of hierarchical classification levels and non-hierarchical categories, and the labels shall be used as the basis for mandatory access control decisions. The TCB shall be able to support two or more such sensitivity levels. (See the Mandatory Access Control interpretations.) The following requirements shall hold for all accesses between all subjects external to the TCB and all objects directly or indirectly accessible by these subjects. A subject can read an object only if the hierarchical classification in the subject's sensitivity level is greater than or equal to the hierarchical classification of the object's sensitivity level and the non-hierarchical categories in the subject's sensitivity level include all the non-hierarchical categories in the object's sensitivity level. A subject can write an object only if the hierarchical classification in the subject's sensitivity level is less than or equal to the hierarchical classification of the object's sensitivity level and the non-hierarchical categories in the subject's sensitivity level are included in the non-hierarchical categories in the object's sensitivity level. Identification and authentication data shall be used by the TCB to authenticate the user's identity and to ensure that the sensitivity level and authorization of subjects external to the TCB that may be created to act on behalf of the individual user are dominated by the clearance and authorization of that user.

+ Interpretation

Each partition of the NTCB exercises mandatory access control policy over all subjects and objects in its component. In a network, the responsibility of an NTCB partition encompasses all mandatory access control functions in its component that would be required of a TCB in a stand-alone system. In particular, subjects and objects used for communication with other components are under the control of the NTCB partition. Mandatory access control includes secrecy and integrity control to the extent that the network sponsor has described in the overall network security policy.

Conceptual entities associated with communication between two components, such as sessions, connections and virtual circuits, may be thought of as having two ends, one in each component, where each end is represented by a local object. Communication is viewed as an operation that copies information from an object at one end of a communication path to an object at the other end. Transient data-carrying entities, such as datagrams and packets, exist either as information within other objects, or as a pair of objects, one at each end of the communication path.

The requirement for "two or more" sensitivity levels can be met by either secrecy or integrity levels. When there is a mandatory integrity policy, the stated requirements for reading and writing are generalized to: A subject can read an object only if the subject's sensitivity level dominates the object's sensitivity level, and a subject can write an object only if the object's sensitivity level dominates the subject's sensitivity level. Based on the integrity policy, the network sponsor shall define the dominance relation for the total label, for example, by combining secrecy and integrity lattices. -

+ Rationale

An NTCB partition can maintain access control only over subjects and objects in its component. At levels B2 and above, the NTCB partition must maintain access control over all subjects and objects in its component. Access by a subject in one component to information contained in an object in another component requires the creation of a subject in the remote component which acts as a surrogate for the first subject.

The mandatory access controls must be enforced at the interface of the reference monitor (viz. the mechanism that controls physical processing resources) for each NTCB partition. This mechanism creates the abstraction of subjects and objects which it controls. Some of these subjects outside the reference monitor, per se, may be designated to implement part of an NTCB partition's mandatory policy,

e.g., by using the "trusted subjects" defined in the Bell-LaPadula model.

The prior requirements on exportation of labeled information to and from I/O devices ensure the consistency between the sensitivity labels of objects connected by a communication path. As noted in the introduction, the network architecture must recognize the linkage between the overall mandatory network security policy and the connection oriented abstraction. For example, individual data-carrying entities such as datagrams can have individual sensitivity labels that subject them to mandatory access control in each component. The abstraction of a single-level connection is realized and enforced implicitly by an architecture while a connection is realized by single-level subjects that necessarily employ only datagrams of the same level.

The fundamental trusted systems technology permits the DAC mechanism to be distributed, in contrast to the requirements for mandatory access control. For networks this separation of MAC and DAC mechanisms is the rule rather than

- See, for example, Grohn, M. J., A Model of a Protected Data Management System, ESD-TR-76-289, I. P. Sharp Assoc. Ltd., June, 1976; and Denning, D. E., Lunt, T. F., Neumann, P. G., Schell, R. R., Heckman, M. and Shockley, W., Secure Distributed Data Views, Security Policy and Interpretation for a Class A1 Multilevel Secure Relational Database System, SRI International, November 1986.

the exception.

The set of total sensitivity labels used to represent all the sensitivity levels for the mandatory access control (combined data secrecy and data integrity) policy always forms a partially ordered set. Without loss of generality, this set of labels can always be extended to form a lattice, by including all the combinations of non-hierarchical categories. As for any lattice, a dominance relation is always defined for the total sensitivity labels. For administrative reasons it may be helpful to have a maximum level which dominates all others.

4.1.2 Accountability

4.1.2.1 Identification and Authentication

+ Statement from DoD 5200.28-STD

The TCB shall require users to identify themselves to it before beginning to perform any other actions that the TCB is expected to mediate. Furthermore, the TCB shall maintain authentication data that includes information for verifying the identify of individual users (e.g., passwords) as well as information for determining the clearance and authorizations of individual users. This data shall be used by the TCB to authenticate the user's identity and to ensure that the sensitivity level and authorization of subjects external to the TCB that may be created to act on behalf of the individual user are dominated by the clearance and authorization of that user. The TCB shall protect authentication data so that it cannot be accessed by any unauthorized user. The TCB shall be able to enforce individual accountability by providing the capability to uniquely identify each individual ADP system user. The TCB shall also provide the capability of associating this identity with all auditable actions taken by that individual.

+ Interpretation

The requirement for identification and authentication of users is the same for a network system as for an ADP system. The identification and authentication may be done by the component to which the user is directly connected or some other component, such as an identification and authentication server. Available techniques, such as those described in the Password Guideline=, are generally also applicable in the network context. However, in cases where the NTCB is expected to mediate actions of a host (or other network component) that is acting on behalf of a user or group of users, the NTCB may employ identification and

= Department of Defense Password Management Guide-
line, CSC-STD-002-85

authentication of the host (or other component) in lieu of identification and authentication of an individual user, so long as the component identifier implies a list of specific users uniquely associated with the identifier at the time of its use for authentication. This requirement does not apply to internal subjects.

Authentication information, including the identity of a user (once authenticated) may be passed from one component to another without reauthentication, so long as the NTCB

protects (e.g., by encryption) the information from unauthorized disclosure and modification. This protection shall provide at least a similar level of assurance (or strength of mechanism) as pertains to the protection of the authentication mechanism and authentication data.

+ Rationale

The need for accountability is not changed in the context of a network system. The fact that the NTCB is partitioned over a set of components neither reduces the need nor imposes new requirements. That is, individual accountability is still the objective. Also, in the context of a network system at the (C2) level or higher "individual accountability" can be satisfied by identification of a host (or other component) so long as the requirement for traceability to individual users or a set of specific individual users with active subjects is satisfied. There is allowed to be an uncertainty in traceability because of elapsed time between changes in the group membership and the enforcement in the access control mechanisms. In addition, there is no need in a distributed processing system like a network to reauthenticate a user at each point in the network where a projection of a user (via the subject operating on behalf of the user) into another remote subject takes place.

The passing of identifiers and/or authentication information from one component to another is usually done in support to the implementation of the discretionary access control (DAC). This support relates directly to the DAC regarding access by a user to a storage object in a different NTCB partition than the one where the user was authenticated. Employing a forwarded identification implies additional reliance on the source and components along the path. If the authenticated identification is used as the basis of determining a sensitivity label for a subject, it must satisfy the Label Integrity criterion.

An authenticated identification may be forwarded between components and employed in some component to identify the sensitivity level associated with a subject created to act on behalf of the user so identified.

4.1.2.1.1 Trusted Path

+ Statement from DoD 5200.28-STD

The TCB shall support a trusted communication path between itself and users for use when a positive TCB-to-user connection is required (e.g., login, change subject sensitivity level). Communications via this trusted path shall be activated exclusively by a user or the TCB and shall be logically and unmistakably distinguishable from other paths.

+ Interpretation

A trusted path is supported between a user (i.e., human) and the NTCB partition in the component to which the user is directly connected.

+ Rationale

When a user logs into a remote component, the user id is transmitted securely between the local and remote NTCB partitions in accordance with the requirements in Identification and Authentication.

Trusted Path is necessary in order to assure that the user is communicating with the NTCB and only the NTCB when security relevant activities are taking place (e.g., authenticate user, set current session sensitivity level). However, Trusted Path does not address communications within the NTCB, only communications between the user and the NTCB. If, therefore, a component does not support any direct user communication then the component need not contain mechanisms for assuring direct NTCB to user communications.

The requirement for trusted communication between one NTCB partition and another NTCB partition is addressed in the System Architecture section. These requirements are separate and distinct from the user to NTCB communication requirement of a trusted path. However, it is expected that this trusted communication between one NTCB partition and another NTCB partition will be used in conjunction with the trusted path to implement trusted communication between the user and the remote NTCB partition.

4.1.2.2 Audit

+ Statement from DoD 5200.28-STD

The TCB shall be able to create, maintain, and protect from modification or unauthorized access or destruction an audit trail of accesses to the objects it protects. The audit data shall be protected by the TCB so that read access to it is limited to those who are authorized for audit data. The TCB shall be able to record the following types of events: use of identification and authentication mechanisms, introduction of objects into a user's address space (e.g., file open, program initiation), deletion of objects, actions taken by computer operators and system administrators and/or system security officers, and other security relevant events. The TCB shall also be able to audit any override of human-readable output markings. For each recorded event, the audit record shall identify: date and time of the event, user, type of event, and success or failure of the event.

For identification/authentication events the origin of request (e.g., terminal ID) shall be included in the audit record. For events that introduce an object into a user's address space and for object deletion events the audit record shall include the name of the object and the object's sensitivity level. The ADP system administrator shall be able to selectively audit the actions of any one or more users based on individual identify and/or object sensitivity level. The TCB shall be able to audit the identified events that may be used in the exploitation of covert storage channels. The TCB shall contain a mechanism that is able to monitor the occurrence or accumulation of security auditable events that may indicate an imminent violation of security policy. This mechanism shall be able to immediately notify the security administrator when thresholds are exceeded and, if the occurrence or accumulation of these security relevant events continues, the system shall take the least disruptive action to terminate the event.

+ Interpretation

This criterion applies as stated. The sponsor must select which events are auditable. If any such events are not distinguishable by the NTCB alone (for example those identified in Part II), the audit mechanism shall provide an interface, which an authorized subject can invoke with parameters sufficient to produce an audit record. These audit records shall be distinguishable from those provided by the NTCB. In the context of a network system, "other security relevant events" (depending on network system architecture and network security policy) might be as follows:

1. Identification of each access event (e.g., establishing a connection or a connectionless association between processes in two hosts of the network) and its principal parameters (e.g., host identifiers of the two hosts involved in the access event and user identifier or host identifier of the user or host that is requesting the access event)
2. Identification of the starting and ending times of each access event using local time or global synchronized time
3. Identification of security-relevant exceptional conditions (e.g., potential violation of data integrity, such as misrouted datagrams) detected during the transactions between two hosts
4. Utilization of cryptographic variables
5. Changing the configuration of the network (e.g., a

component leaving the network and rejoining)

In addition, identification information should be included in appropriate audit trail records, as necessary, to allow association of all related (e.g., involving the same network event) audit trail records (e.g., at different hosts) with each other. Furthermore, a component of the network system may provide the required audit capability (e.g., storage, retrieval, reduction, analysis) for other components that do not internally store audit data but transmit the audit data to some designated collection component. Provisions shall be made to control the loss of audit data due to unavailability of resources.

In the context of a network system, the "user's address space" is extended, for object introduction and deletion events, to include address spaces being employed on behalf of a remote user (or host). However, the focus remains on users in contrast to internal subjects as discussed in the DAC criterion. In addition, audit information must be stored in machine-readable form.

The capability must exist to audit the identified events that may be used in the exploitation of covert storage channels. To accomplish this, each NTCB partition must be able to audit those events locally that may lead to the exploitation of a covert storage channel which exist because of the network.

The sponsor shall identify the specific auditable events that may indicate an imminent violation of security policy. The component which detects the occurrence or accumulation of such events must be able to notify an appropriate administrator when thresholds are exceeded, and to initiate actions which will result in termination of the event if the accumulation continues. For example, when the threshold of unsuccessful login attempts within a period of time is exceeded, login shall be inhibited for a specific time period.

+ Rationale

For remote users, the network identifiers (e.g., internet address) can be used as identifiers of groups of individual users (e.g., "all users at Host A") to eliminate the maintenance that would be required if individual identification of remote users was employed. In this class (C2), however, it must be possible to identify (immediately or at some later time) the individuals represented by a group identifier. In all other respects, the interpretation is a straightforward extension of the criterion into the context of a network system. Identification of covert channel events is addressed in the Covert Channel Analysis section.

Because of concurrency and synchronization problems, it may not be possible to detect in real time the accumulation of security auditable events that are occurring in different NTCB partitions. However, each NTCB partition that has been allocated audit responsibility must have the capability to detect the local accumulation of events, to notify the partition security administrator and/or the network security administrator, and to initiate actions which will result in termination of the event locally.

4.1.3 Assurance

4.1.3.1 Operational Assurance

4.1.3.1.1 System Architecture

+ Statement from DoD 5200.28-STD

The TCB shall maintain a domain for its own execution that protects it from external interference or tampering (e.g., by modification of its code or data structures). The TCB shall maintain process isolation through the provision of distinct address spaces under its control. The TCB shall be internally structured into well-defined largely independent modules. It shall make effective use of available hardware to separate those elements that are protection-critical from those that are not. The TCB modules shall be designed such that the principle of least privilege is enforced. Features in hardware, such as segmentation, shall be used to support logically distinct storage objects with separate attributes (namely: readable, writable). The user interface to the TCB shall be completely defined and all elements of the TCB identified. The TCB shall be designed and structured to use a complete, conceptually simple protection mechanism with precisely defined semantics. This mechanism shall play a central role in enforcing the internal structuring of the TCB and the system. The TCB shall incorporate significant use of layering, abstraction and data hiding. Significant system engineering shall be directed toward minimizing the complexity of the TCB and excluding from the TCB modules that are not protection-critical.

+ Interpretation

The system architecture criterion must be met individually by all NTCB partitions. Implementation of the requirement that the NTCB maintain a domain for its own execution is achieved by having each NTCB partition maintain a domain for its own execution. Since each component is itself a distinct domain in the overall network system, this also satis-

fies the requirement for process isolation through distinct address spaces in the special case where a component has only a single subject.

The NTCB must be internally structured into well-defined largely independent modules and meet the hardware requirements. This is satisfied by having each NTCB partition so structured. The NTCB controls all network resources. These resources are the union of the sets of resources over which the NTCB partitions have control. Code and data structures belonging to the NTCB, transferred among NTCB subjects (i.e., subjects outside the reference monitor but inside the NTCB) belonging to different NTCB partitions, must be protected against external interference or tampering. For example, a cryptographic checksum or physical means may be employed to protect user authentication data exchanged between NTCB partitions.

Each NTCB partition must enforce the principle of least privilege within its component. Additionally, the NTCB must be structured so that the principle of least privilege is enforced in the system as a whole.

The NTCB must be designed and structured according to the network security architecture to use a complete, conceptually simple protection mechanism. Furthermore, each NTCB partition must also be so designed and structured.

Significant system engineering should be directed toward minimizing the complexity of each NTCB partition, and of the NTCB. Care shall be taken to exclude modules (and components) that are not protection-critical from the NTCB.

It is recognized that some modules and/or components may need to be included in the NTCB and must meet the NTCB requirements even though they may not appear to be directly protection-critical. The correct operation of these modules/components is necessary for the correct operation of the protection-critical modules and components. However, the number and size of these modules/components should be kept to a minimum.

Each NTCB partition provides isolation of resources (within its component) in accord with the network system architecture and security policy so that "supporting elements" (e.g., DAC and user identification) for the security mechanisms of the network system are strengthened compared to C2, from an assurance point of view, through the provision of distinct address spaces under control of the NTCB.

As discussed in the Discretionary Access Control section, the DAC mechanism of a NTCB partition may be implemented at the interface of the reference monitor or may be distributed in subjects that are part of the NTCB in the

same or different component. When distributed in NTCB subjects (i.e., when outside the reference monitor), the assurance requirements for the design and implementation of the DAC shall be those of class C2 for all networks of class C2 or above.

+ Rationale

The requirement that the NTCB be structured into modules and meet the hardware requirements applies within the NTCB partitions in the various components.

The principle of least privilege requires that each user or other individual with access to the system be given only those resources and authorizations required for the performance of this job. In order to enforce this principle in the system it must be enforced in every NTCB partition that supports users or other individuals. For example, prohibiting access by administrators to objects outside the NTCB partition (e.g., games) lessens the opportunity of damage by a Trojan Horse.

The requirement for the protection of communications between NTCB partitions is specifically directed to subjects that are part of the NTCB partitions. Any requirements for such protection for the subjects that are outside the NTCB partitions are addressed in response to the integrity requirements of the security policy.

There are certain parts of a network (modules and/or components) that may not appear to be directly protection-critical in that they are not involved in access control decisions, do not directly audit, and are not involved in the identification/authentication process. However, the security of the network must depend on the correct operation of these modules and/or components. An example of this is a single level packet switch. Although it may not normally be involved directly in enforcing the discretionary security policy, this switch may be trusted not to mix data from different message streams. If the switch does not operate correctly, data could get mixed, and unauthorized access could result. Therefore, these modules/components must be included in the NTCB and must meet the NTCB requirements applicable to the policy element(s) for which they are responsible.

4.1.3.1.2 System Integrity

+ Statement from DoD 5200.28-STD

Hardware and/or software features shall be provided that can be used to periodically validate the correct operation of

the on-site hardware and firmware elements of the TCB.

+ Interpretation

Implementation of the requirement is partly achieved by having hardware and/or software features that can be used to periodically validate the correct operation of the hardware and firmware elements of each component's NTCB partition. Features shall also be provided to validate the identity and correct operation of a component prior to its incorporation in the network system and throughout system operation. For example, a protocol could be designed that enables the components of the partitioned NTCB to exchange messages periodically and validate each other's correct response. The protocol shall be able to determine the remote entity's ability to respond. NTCB partitions shall provide the capability to report to network administrative personnel the failures detected in other NTCB partitions.

Intercomponent protocols implemented within a NTCB shall be designed in such a way as to provide correct operation in the case of failures of network communications or individual components. The allocation of mandatory and discretionary access control policy in a network may require communication between trusted subjects that are part of the NTCB partitions in different components. This communication is normally implemented with a protocol between the subjects as peer entities. Incorrect access within a component shall not result from failure of an NTCB partition to communicate with other components.

+ Rationale

The first paragraph of the interpretation is a straightforward extension of the requirement into the context of a network system and partitioned NTCB as defined for these network criteria.

NTCB protocols should be robust enough so that they permit the system to operate correctly in the case of localized failure. The purpose of this protection is to preserve the integrity of the NTCB itself. It is not unusual for one or more components in a network to be inoperative at any time, so it is important to minimize the effects of such failures on the rest of the network. Additional integrity and denial of service issues are addressed in Part II.

It should be clear that some integrity and denial of service features can reside outside the NTCB. Otherwise all software in a network would be in the NTCB. Every piece of software that has an opportunity to write to some data or protocol field is "trusted" to preserve integrity or not cause denial of service to some extent. For example, it is necessary to "trust" TELNET to correctly translate user

data, and to eventually transmit packets. FTP also has to be "trusted" to not inappropriately modify files, and to attempt to complete the file transfer. These protocols can be designed, however to exist outside the NTCB (from a protection perspective). It is beneficial to do this type of security engineering so that the amount of code that must be trusted to not disclose data is minimized. Putting everything inside the NTCB contradicts the requirement to perform "significant system engineering ... directed toward ... excluding from the TCB modules that are not protection critical," which removes the primary difference between B2 and B3. If everything has to be in the TCB to ensure data integrity and protection against denial of service, there will be considerably less assurance that disclosure protection is maximized.

4.1.3.1.3 Covert Channel Analysis

+ Statement from DoD 5200.28-STD

The system developer shall conduct a thorough search for covert channels and make a determination (either by actual measurement or by engineering estimation) of the maximum bandwidth of each identified channel. (See the Covert Channels Guideline section.) FORMAL METHODS SHALL BE USED IN THE ANALYSIS.

+ Interpretation

The requirement, including the TCSEC Covert Channel Guideline, applies as written. In a network, there are additional instances of covert channels associated with communication between components. THE FORMAL METHODS SHALL BE USED IN THE ANALYSIS OF EACH INDIVIDUAL COMPONENT DESIGN AND IMPLEMENTATION.

+ Rationale

The exploitation of network protocol information (e.g., headers) can result in covert storage channels. Exploitation of frequency of transmission can result in covert timing channels. The topic has been addressed in the literature.-

4.1.3.1.4 Trusted Facility Management

+ Statement from DoD 5200.28-STD

The TCB shall support separate operator and administrator functions. The functions performed in the role of a security administrator shall be identified. The ADP system

administrative personnel shall only be able to perform security administrator functions after taking a distinct auditable action to assume the security administrator role on the ADP system. Non-security functions that can be performed in the security administration role shall be limited strictly

- See, for example, Girling, C. G., "Covert Channels in LAN's," IEEE Transactions on Software Engineering, Vol. SE-13, No. 2, February 1987; and Padlipsky, M. A., Snow, D. P., and Karger, P. A., Limitations of End-to-End Encryption in Secure Computer Networks, MITRE Technical Report, MTR-3592, Vol. I, May 1978 (ESD TR 78-158, DTIC AD A059221).

to those essential to performing the security role effectively.

+ Interpretation

This requirement applies as written to both the network as a whole and to individual components which support such personnel.

+ Rationale

It is recognized that based on the allocated policy elements some components may operate with no human interface.

4.1.3.1.5 Trusted Recovery

+ Statement from DoD 5200.28-STD

Procedures and/or mechanisms shall be provided to assure that, after an ADP system failure or other discontinuity, recovery without a protection compromise is obtained.

+ Interpretation

The recovery process must be accomplished without a protection compromise after the failure or other discontinuity of any NTCB partition. It must also be accomplished after a failure of the entire NTCB.

+ Rationale

This is a straight-forward extension of the requirement into the network context, and takes into account that it is possible for parts of the system to fail while other parts

continue to operate normally. This may be a security-relevant event; if so it must be audited.

4.1.3.2 Life-Cycle Assurance

4.1.3.2.1 Security Testing

+ Statement from DoD 5200.28-STD

The security mechanisms of the ADP system shall be tested and found to work as claimed in the system documentation. A team of individuals who thoroughly understand the specific implementation of the TCB shall subject its design documentation, source code, and object code to thorough analysis and testing. Their objectives shall be: to uncover all design and implementation flaws that would permit a subject external to the TCB to read, change, or delete data normally denied under the mandatory or discretionary security policy enforced by the TCB; as well as to assure that no subject (without authorization to do so) is able to cause the TCB to enter a state such that it is unable to respond to communications initiated by other users. The TCB shall be found resistant to penetration. All discovered flaws shall be removed or neutralized and the TCB retested to demonstrate that they have been eliminated and that new flaws have not been introduced. Testing shall demonstrate that the TCB implementation is consistent with the descriptive top-level specification. No design flaws and no more than a few correctable implementation flaws may be found during testing and there shall be reasonable confidence that few remain. MANUAL OR OTHER MAPPING OF THE FTLS TO THE SOURCE CODE MAY FORM A BASIS FOR PENETRATION TESTING. (See the Security Testing Guidelines.)

+ Interpretation

Testing of a component will require a testbed that exercises the interfaces and protocols of the component including tests under exceptional conditions. The testing of a security mechanism of the network system for meeting this criterion shall be an integrated testing procedure involving all components containing an NTCB partition that implement the given mechanism. This integrated testing is additional to any individual component tests involved in the evaluation of the network system. The sponsor should identify the allowable set of configurations including the sizes of the networks. Analysis or testing procedures and tools shall be available to test the limits of these configurations. A change in configuration within the allowable set of configurations does not require retesting.

The testing of each component will include the intro-

duction of subjects external to the NTCB partition for the component that will attempt to read, change, or delete data normally denied. If the normal interface to the component does not provide a means to create the subjects needed to conduct such a test, then this portion of the testing shall use a special version of the untrusted software for the component that results in subjects that make such attempts. The results shall be saved for test analysis. Such special versions shall have an NTCB partition that is identical to that for the normal configuration of the component under evaluation.

The testing of the mandatory controls shall include tests to demonstrate that the labels for information imported and/or exported to/from the component accurately represent the labels maintained by the NTCB partition for the component for use as the basis for its mandatory access control decisions. The tests shall include each type of device, whether single-level or multilevel, supported by the component.

The NTCB must be found resistant to penetration. This applies to the NTCB as a whole, and to each NTCB partition in a component of this class.

+ Rationale

The phrase "no subject (without authorization to do so) is able to cause the TCB to enter a state such that it is unable to respond to communications initiated by other users" relates to the security services (Part II of this TNI) for the Denial of Service problem, and to correctness of the protocol implementations.

Testing is an important method available in this evaluation division to gain any assurance that the security mechanisms perform their intended function. A major purpose of testing is to demonstrate the system's response to inputs to the NTCB partition from untrusted (and possibly malicious) subjects.

In contrast to general purpose systems that allow for the dynamic creation of new programs and the introductions of new processes (and hence new subjects) with user specified security properties, many network components have no method for introducing new programs and/or processes during their normal operation. Therefore, the programs necessary for the testing must be introduced as special versions of the software rather than as the result of normal inputs by the test team. However, it must be insured that the NTCB partition used for such tests is identical to the one under evaluation.

Sensitivity labels serve a critical role in maintaining

the security of the mandatory access controls in the network. Especially important to network security is the role of the labels for information communicated between components - explicit labels for multilevel devices and implicit labels for single-level devices. Therefore the testing for correct labels is highlighted.

The requirement for testing to demonstrate consistency between the NTCB implementation and the FTLS is a straightforward extension of the TCSEC requirement into the context of a network system.

4.1.3.2.2 Design Specification and Verification

+ Statement from DoD 5200.28-STD

A formal model of the security policy supported by the TCB shall be maintained over the life cycle of the ADP system that is proven and demonstrated to be consistent with its axioms. A descriptive top-level specification (DTLS) of the TCB shall be maintained that completely and accurately describes the TCB in terms of exceptions, error messages, and effects. A FORMAL TOP-LEVEL SPECIFICATION (FTLS) OF THE TCB SHALL BE MAINTAINED THAT ACCURATELY DESCRIBES THE TCB IN TERMS OF EXCEPTIONS, ERROR MESSAGES, AND EFFECTS. THE DTLS AND FTLS SHALL INCLUDE THOSE COMPONENTS OF THE TCB THAT ARE IMPLEMENTED AS HARDWARE AND/OR FIRMWARE IF THEIR PROPERTIES ARE VISIBLE AT THE TCB INTERFACE. THE FTLS SHALL BE SHOWN to be an accurate description of the TCB interface. A convincing argument shall be given that the DTLS is consistent with the model AND A COMBINATION OF FORMAL AND INFORMAL TECHNIQUES SHALL BE USED TO SHOW THAT THE FTLS IS CONSISTENT WITH THE MODEL. THIS VERIFICATION EVIDENCE SHALL BE CONSISTENT WITH THAT PROVIDED WITHIN THE STATE-OF-THE-ART OF THE PARTICULAR NATIONAL COMPUTER SECURITY CENTER-ENDORSED FORMAL SPECIFICATION AND VERIFICATION SYSTEM USED. MANUAL OR OTHER MAPPING OF THE FTLS TO THE TCB SOURCE CODE SHALL BE PERFORMED TO PROVIDE EVIDENCE OF CORRECT IMPLEMENTATION.

+ Interpretation

The overall network security policy expressed in this model will provide the basis for the mandatory access control policy exercised by the NTCB over subjects and storage objects in the entire network. The policy will also be the basis for the discretionary access control policy exercised by the NTCB to control access of named users to named objects. Data integrity requirements addressing the effects of unauthorized MSM need not be included in this model. The overall network policy must be decomposed into policy elements that are allocated to appropriate components and used as the basis for the security policy model for those com-

ponents.

The level of abstraction of the model, and the set of subjects and objects that are explicitly represented in the model, will be affected by the NTCB partitioning. Subjects and objects must be represented explicitly in the model for the partition if there is some network component whose NTCB partition exercises access control over them. The model shall be structured so that the axioms and entities applicable to individual network components are manifest. Global network policy elements that are allocated to components shall be represented by the model for that component.

AN FTLS FOR A NETWORK CONSISTS OF A COMPONENT FTLS FOR EACH UNIQUE TRUSTED NETWORK COMPONENT, PLUS ANY GLOBAL DECLARATIONS AND ASSERTIONS THAT APPLY TO MORE THAN ONE COMPONENT. IF THE MODEL FOR EACH COMPONENT REPRESENTS ALL THE GLOBAL MANDATORY POLICY ELEMENTS ALLOCATED TO THAT COMPONENT, THERE MAY NOT BE ANY GLOBAL ASSERTIONS NEEDED, AND IN THIS CASE THE COLLECTION OF COMPONENT FTLS, WITH ANY SHARED DECLARATIONS, IS THE NETWORK FTLS. EACH COMPONENT FTLS SHALL DESCRIBE THE INTERFACE TO THE NTCB PARTITION OF ITS COMPONENTS. The requirements for a network DTLs are given in the Design Documentation section.

+ Rationale

The treatment of the model depends to a great extent on the degree of integration of the communications service into a distributed system. In a closely coupled distributed system, one might use a model that closely resembles one appropriate for a stand-alone computer system.

In all cases, the model of each partition will be expected to show the role of the NTCB partition in each kind of component. It will most likely clarify the model, although not part of the model, to show access restrictions implied by the system design; for example, subjects representing protocol entities might have access only to objects containing data units at the same layer of protocol. The allocation of subjects and objects to different protocol layers is a protocol design choice which need not be reflected in the security policy model.

THE FTLS MUST REPRESENT THE UNDERLYING REFERENCE MONITOR AND ANY SUBJECTS IMPLEMENTING THE MANDATORY POLICY. OTHER POLICY ELEMENTS DISTRIBUTED IN NTCB SUBJECTS (SEE THE INTERPRETATION OF SYSTEM ARCHITECTURE) NEED NOT BE REPRESENTED BY THE FTLS.

4.1.3.2.3 Configuration Management

+ Statement from DoD 5200.28-STD

During THE ENTIRE LIFE-CYCLE, I.E. DURING THE DESIGN, DEVELOPMENT, and maintenance of the TCB, a configuration management system shall be in place FOR ALL SECURITY-RELEVANT HARDWARE, FIRMWARE, AND SOFTWARE that maintains control of changes to THE FORMAL MODEL, the descriptive AND FORMAL top-level SPECIFICATIONS, other design data, implementation documentation, source code, the running version of the object code, and test fixtures and documentation. The configuration management system shall assure a consistent mapping among all documentation and code associated with the current version of the TCB. Tools shall be provided for generation of a new version of the TCB from source code. Also available shall be tools, MAINTAINED UNDER STRICT CONFIGURATION CONTROL, for comparing a newly generated version with the previous TCB version in order to ascertain that only the intended changes have been made in the code that will actually be used as the new version of the TCB. A COMBINATION OF TECHNICAL, PHYSICAL, AND PROCEDURAL SAFEGUARDS SHALL BE USED TO PROTECT FROM UNAUTHORIZED MODIFICATION OR DESTRUCTION THE MASTER COPY OR COPIES OF ALL MATERIAL USED TO GENERATE THE TCB.

+ Interpretation

The requirement applies as written, with the following extensions:

1. A configuration management system must be in place for each NTCB partition.
2. A configuration management plan must exist for the entire system. If the configuration management system is made up of the conglomeration of the configuration management systems of the various NTCB partitions, then the configuration management plan must address the issue of how configuration control is applied to the system as a whole.

ALL MATERIAL USED IN GENERATING A NEW VERSION OF THE NTCB AND EACH NTCB PARTITION MUST BE PROTECTED, REGARDLESS OF WHERE IT PHYSICALLY RESIDES.

+ Rationale

Each NTCB partition must have a configuration management system in place, or else there will be no way for the NTCB as a whole to have an effective configuration management system. The other extensions are merely reflections of the way that networks operate in practice.

THIS NEW REQUIREMENT EXPLICITLY MANDATES THE PROTECTION OF MATERIAL USED TO GENERATE AN NTCB PARTITION, EVEN WHEN

THE GENERATION OCCURS BY DOWN-LINE LOADING OF A REMOTE COMPONENT.

4.1.3.2.4 Trusted Distribution

+ Statement from DoD 5200.28-STD

A TRUSTED ADP SYSTEM CONTROL AND DISTRIBUTION FACILITY SHALL BE PROVIDED FOR MAINTAINING THE INTEGRITY OF THE MAPPING BETWEEN THE MASTER DATA DESCRIBING THE CURRENT VERSION OF THE TCB AND THE ON-SITE MASTER COPY OF THE CODE FOR THE CURRENT VERSION. PROCEDURES (E.G., SITE SECURITY ACCEPTANCE TESTING) SHALL EXIST FOR ASSURING THAT THE TCB SOFTWARE, FIRMWARE, AND HARDWARE UPDATES DISTRIBUTED TO A CUSTOMER ARE EXACTLY AS SPECIFIED BY THE MASTER COPIES.

+ Interpretation

THIS REQUIREMENT APPLIES AS STATED, WITH THE ADDITIONAL REQUIREMENT THAT, IF DOWN-LINE LOADING IS USED, THERE MUST BE A TRUSTED METHOD OF GENERATING, SENDING, AND LOADING ANY SOFTWARE INVOLVED.

+ Rationale

THIS IS A STRAIGHTFORWARD EXTENSION OF THE REQUIREMENT INTO THE NETWORK CONTEXT.

4.1.4 Documentation.

4.1.4.1 Security Features User's Guide

+ Statement from DoD 5200.28-STD

A single summary, chapter, or manual in user documentation shall describe the protection mechanisms provided by the TCB, interpretations on their use, and how they interact with one another.

+ Interpretation

This user documentation describes user visible protection mechanisms at the global (network system) level and at the user interface of each component, and the interaction among these.

+ Rationale

The interpretation is an extension of the requirement into the context of a network system as defined for these network criteria. Documentation of protection mechanisms provided by individual components is required by the criteria for trusted computer systems that are applied as appropriate for the individual components.

4.1.4.2 Trusted Facility Manual

+ Statement from DoD 5200.28-STD

A manual addressed to the ADP system administrator shall present cautions about functions and privileges that should be controlled when running a secure facility. The procedures for examining and maintaining the audit files as well as the detailed audit record structure for each type of audit event shall be given. The manual shall describe the operator and administrator functions related to security, to include changing the security characteristics of a user. It shall provide interpretations on the consistent and effective use of the protection features of the system, how they interact, how to securely generate a new TCB, and facility procedures, warnings, and privileges that need to be controlled in order to operate the facility in a secure manner. The TCB modules that contain the reference validation mechanism shall be identified. The procedures for secure generation of a new TCB from source after modification of any modules in the TCB shall be described. It shall include the procedures to ensure that the system is initially started in a secure manner. Procedures shall also be included to resume secure system operation after any lapse in system operation.

+ Interpretation

This manual shall contain specifications and procedures to assist the system administrator(s) maintain cognizance of the network configuration. These specifications and procedures shall address the following:

1. The hardware configuration of the network itself;
2. The implications of attaching new components to the network;
3. The case where certain components may periodically leave the network (e.g., by crashing, or by being disconnected) and then rejoin;
4. Network configuration aspects that can impact the security of the network system; (For example, the manual should describe for the network system administrator the interconnections among components that are consistent with the overall network system

architecture.)

5. Loading or modifying NTCB software or firmware (e.g., down-line loading).
6. Incremental updates; that is, it must explicitly indicate which components of the network may change without others also changing.

The physical and administrative environmental controls shall be specified. Any assumptions about security of a given network should be clearly stated (e.g., the fact that all communications links must be physically protected to a certain level).

The components of the network that form the NTCB must be identified. Furthermore, the modules within an NTCB partition that contain the reference validation mechanism (if any) within that partition must be identified.

The procedures for the secure generation of a new version (or copy) of each NTCB partition from source must be described. The procedures and requirements for the secure generation of the NTCB necessitated by changes in the network configuration shall be described.

Procedures for starting each NTCB partition in a secure state shall be specified. Procedures must also be included to resume secure operation of each NTCB partition and/or the NTCB after any lapse in system or subsystem operation.

+ Rationale

There may be multiple system administrators with diverse responsibilities. The technical security measures described by these criteria must be used in conjunction with other forms of security in order to achieve security of the network. Additional forms include administrative security, physical security, emanations security, etc.

Extension of this criterion to cover configuration aspects of the network is needed because, for example, proper interconnection of components is typically essential to achieve a correct realization of the network architecture.

As mentioned in the section on Label Integrity, cryptography is one common mechanism employed to protect communication circuits. Encryption transforms the representation of information so that it is unintelligible to unauthorized subjects. Reflecting this transformation, the sensitivity of the ciphertext is generally lower than the cleartext. If encryption methodologies are employed, they shall be approved by the National Security Agency (NSA).

The encryption algorithm and its implementation are outside the scope of these interpretations. This algorithm and implementation may be implemented in a separate device or may be a function of a subject in a component not dedicated to encryption. Without prejudice, either implementation packaging is referred to as an encryption mechanism herein.

The requirements for descriptions of NTCB generation and identification of modules and components that form the NTCB are straightforward extensions of the TCSEC requirements into the network context. In those cases where the vendor does not provide source code, an acceptable procedure shall be to request the vendor to perform the secure generation.

Given the nature of network systems (e.g., various components tend to be down at different times, and the network system must continue operation without that component), it is imperative to know both how to securely start up an NTCB partition, and how to resume operation securely. It is also necessary to know how to resume secure operation of the NTCB after any partition has been down.

4.1.4.3 Test Documentation

+ Statement from DoD 5200.28-STD

The system developer shall provide to the evaluators a document that describes the test plan, test procedures that show how the security mechanisms were tested, and results of the security mechanisms' functional testing. It shall include results of testing the effectiveness of the methods used to reduce covert channel bandwidths. THE RESULTS OF THE MAPPING BETWEEN THE FORMAL TOP-LEVEL SPECIFICATION AND THE TCB SOURCE CODE SHALL BE GIVEN.

+ Interpretation

The "system developer" is interpreted as "the network system sponsor". The description of the test plan should establish the context in which the testing was or should be conducted. The description should identify any additional test components that are not part of the system being evaluated. This includes a description of the test-relevant functions of such test components and a description of the interfacing of those test components to the system being evaluated. The description of the test plan should also demonstrate that the tests adequately cover the network security policy. The tests should include the features described in the System Architecture and the System

Integrity sections. The tests should also include network configuration and sizing.

THE MAPPING BETWEEN THE FTLS AND THE NTCB SOURCE CODE MUST BE CHECKED TO ENSURE TO THE EXTENT POSSIBLE THAT THE FTLS IS A CORRECT REPRESENTATION OF THE SOURCE CODE, AND THAT THE FTLS HAS BEEN STRICTLY ADHERED TO DURING THE DESIGN AND DEVELOPMENT OF THE NETWORK SYSTEM. THIS CHECK MUST BE DONE FOR EACH COMPONENT OF THE NETWORK SYSTEM FOR WHICH AN FTLS EXISTS.

+ Rationale

The entity being evaluated may be a networking subsystem (see Appendix A) to which other components must be added to make a complete network system. In that case, this interpretation is extended to include contextual definition because, at evaluation time, it is not possible to validate the test plans without the description of the context for testing the networking subsystem.

The bandwidths of covert channels are used to determine the suitability of a network system for a given environment. The effectiveness of the methods used to reduce these bandwidths must therefore be accurately determined.

4.1.4.4 Design Documentation

+ Statement from DoD 5200.28-STD

Documentation shall be available that provides a description of the manufacturer's philosophy of protection and an explanation of how this philosophy is translated into the TCB. The interfaces between the TCB modules shall be described. A formal description of the security policy model enforced by the TCB shall be available and an explanation provided to show that it is sufficient to enforce the security policy. The specific TCB protection mechanisms shall be identified and an explanation given to show that they satisfy the model. The descriptive top-level specification (DTLS) shall be shown to be an accurate description of the TCB interface. Documentation shall describe how the TCB implements the reference monitor concept and give an explanation why it is tamper resistant, cannot be bypassed, and is correctly implemented. The TCB implementation (i.e., in hardware, firmware, and software) shall be informally shown to be consistent with the FORMAL TOP-LEVEL SPECIFICATION (FTLS). The elements of the FTLS shall be shown, using informal techniques, to correspond to the elements of the TCB. Documentation shall describe how the TCB is structured to facilitate testing and to enforce least privilege. This documentation shall also present the results of the covert channel analysis and the tradeoffs involved in restricting the chan-

nels. All auditable events that may be used in the exploitation of known covert storage channels shall be identified. The bandwidths of known covert storage channels, the use of which is not detectable by the auditing mechanisms, shall be provided. (See the Covert Channel Guideline section.) HARDWARE, FIRMWARE, AND SOFTWARE MECHANISMS NOT DEALT WITH IN THE FTLS BUT STRICTLY INTERNAL TO THE TCB (E.G., MAPPING REGISTERS, DIRECT MEMORY ACCESS I/O) SHALL BE CLEARLY DESCRIBED.

+ Interpretation

Explanation of how the sponsor's philosophy of protection is translated into the NTCB shall include a description of how the NTCB is partitioned. The security policy also shall be stated. The description of the interfaces between the NTCB modules shall include the interface(s) between NTCB partitions and modules within the partitions if the modules exist. The sponsor shall describe the security architecture and design, including the allocation of security requirements among components.

The documentation includes both a system description and a set of component DTLs's. The system description addresses the network security architecture and design by specifying the types of components in the network, which ones are trusted, and in what way they must cooperate to support network security objectives. A component DTLs shall be provided for each trusted network component, i.e., each component containing an NTCB partition. Each component DTLs shall describe the interface to the NTCB partition of its component. Both the system description and each component DTLs shall be shown consistent with those assertions in the model that apply to it. Appendix A addresses component evaluation issues.

To show the correspondence between the FTLS and the NTCB implementation, it suffices to show correspondence between each component FTLS and the NTCB partition in that component.

As stated in the introduction to Division B, the sponsor must demonstrate that the NTCB employs the reference monitor concept. The security policy model must be a model for a reference monitor.

The security policy model for each partition implementing a reference monitor shall fully represent the access control policy supported by the partition, including the discretionary and mandatory security policy for secrecy and/or integrity. For the mandatory policy the single dominance relation for sensitivity labels, including secrecy and/or integrity components, shall be precisely defined.

+ Rationale

The interpretation is a straightforward extension of the requirement into the context of a network system as defined for this network interpretation. Other documentation, such as description of components and description of operating environment(s) in which the networking subsystem or network system is designed to function, is required elsewhere, e.g., in the Trusted Facility Manual.

In order to be evaluated, a network must possess a coherent Network Security Architecture and Design. (Interconnection of components that do not adhere to such a single coherent Network Security Architecture is addressed in the Interconnection of Accredited AIS, Appendix C.) The Network Security Architecture must address the security-relevant policies, objectives, and protocols. The Network Security Design specifies the interfaces and services that must be incorporated into the network so that it can be evaluated as a trusted entity. There may be multiple designs that conform to the same architecture but are more or less incompatible and non-interoperable (except through the Interconnection Rules). Security related mechanisms requiring cooperation among components are specified in the design in terms of their visible interfaces; mechanisms having no visible interfaces are not specified in this document but are left as implementation decisions.

The Network Security Architecture and Design must be available from the network sponsor before evaluation of the network, or any component, can be undertaken. The Network Security Architecture and Design must be sufficiently complete, unambiguous, and free from obvious flaws to permit the construction or assembly of a trusted network based on the structure it specifies.

When a component is being designed or presented for evaluation, or when a network assembled from components is assembled or presented for evaluation, there must be a priori evidence that the Network security Architecture and Design are satisfied. That is, the components can be assembled into a network that conforms in every way with the Network Security Architecture and Design to produce a physical realization that is trusted to the extent that its evaluation indicates.

In order for a trusted network to be constructed from components that can be built independently, the Network Security Architecture and Design must completely and unambiguously define the security functionality of components as well as the interfaces between or among components. The Network Security Architecture and Design must be evaluated to determine that a network constructed to its specifications will in fact be trusted, that is, it will be

evaluatable under these interpretations.

The term "model" is used in several different ways in a network context, e.g., a "protocol reference model," a "formal network model," etc. Only the "security policy model" is addressed by this requirement and is specifically intended to model the interface, viz., "security perimeter," of the reference monitor and must meet all the requirements defined in the TCSEC. It must be shown that all parts of the TCB are a valid interpretation of the security policy model, i.e., that there is no change to the secure state except as represented by the model.

Part II: Other Security Services

5. Introduction

Part I of this Interpretation contains interpretations of the Department of Defense Trusted Computer System Evaluation Criteria (TCSEC), DOD 5200.28-STD. Part I deals with controlling access to information. Part II contains additional network security concerns. These concerns differentiate the network environment from the stand-alone computer. Some concerns take on increased significance in the network environment; other concerns do not exist on stand-alone computers. Some of these concerns are outside the scope of Part I; others lack the theoretical basis and formal analysis underlying Part I. The criteria in this Part II address these concerns in the form of additional security requirements that may vary among applications. Overlap between Part I and Part II is minimized as much as possible. However, when an overlap occurs the association between the concerns addressed in both parts is defined. Part II services may be provided by mechanisms outside the NTCB.

5.1. Purpose and Scope

This Part II addresses network security disjoint from Part I. The rating derived in Part I is not effected by Part II. Every component or system must have a Part I evaluation as a basis for the Part II evaluation. Part II includes generic requirements, security features, and evaluation criteria. As described below, Part II evaluations differ from Part I. The purpose of these evaluations is similar, however: to provide guidance to network managers and accreditors as to the reliance they can place in security services. These evaluations are input to the accreditor's decisions concerning the operational mode and range of sensitive

information entrusted to the network.

The network sponsor shall identify the security services offered by his system or component(s). Those services will be evaluated against the criteria for those services in Part II.

5.2. Criteria Form

The general form of Part II criteria is a relatively brief statement, followed by a discussion of functionality, strength of mechanism, and assurance, as appropriate.

Functionality refers to the objective and approach of a security service; it includes features, mechanism, and performance. Alternative approaches to achieving the desired functionality may be more suitable in different applications environments.

Strength of mechanism refers to how well a specific approach may be expected to achieve its objectives. In some cases selection of parameters, such as number of bits used in a checksum or the number of permutations used in an encryption algorithm, can significantly affect strength of mechanism.

Assurance refers to a basis for believing that the functionality will be achieved; it includes tamper resistance, verifiability, and resistance against circumvention or bypass. Assurance is generally based on analysis involving theory, testing, software engineering, validation and verification, and related approaches. The analysis may be formal or informal, theoretical or applied.

For example, consider communications integrity protection against message stream modification. A functionality decision is to select error detection only, or detection and correction; also one may select whether it is sufficient to detect an odd number of bit errors, error bursts of specified duration, or a specified probability of an undetected error. Available mechanisms include parity, longitudinal redundancy check (LRC), cyclic redundancy check (CRC), and cryptographic checkfunction. The strength of the CRC is measured in the probability of an undetected error; this is dependent upon the number of bits employed in the CRC. There is no assurance of security associated with any of the mentioned mechanisms except cryptographic checkfunction. The algorithms are well known; an adversary could change message contents and recalculate the non-cryptographic

checkfunction. The recipient would calculate the checkfunction and not discover that the message had been manipulated. A cryptographic checkfunction would be resistant to such manipulation.

5.3. Evaluation Ratings

Part II evaluations are qualitative, as compared with the hierarchically-ordered ratings (e.g., C1, C2, ...) resulting from Part I. At present it is not considered possible or desirable to employ the same ratings scale in Part II. The results of a Part II evaluation for offered services will generally be summarized using the terms none, minimum, fair, and good. Services not offered by the sponsor will be assigned a rating of not offered. For some services it will be most meaningful to assign a rating of none or present.

The term none is used when the security service is not offered. In some cases the functionality evaluations may be limited to present or none.

The assurance rating for each service is bounded by the Part I or Appendix A evaluation as appropriate because the integrity of the service depends on the protection of the NTCB. Table II-1 relates the Part II assurance rating to the minimum corresponding Part I evaluation ratings.

These Part II evaluations tend to be more qualitative and subjective, and will exhibit greater variance than the Part I evaluations. Nevertheless, Part II evaluations are valuable information concerning the capabilities of the evaluated systems and their suitability for specific applications environments. If functionality, strength of mechanism, and assurance are separately evaluated then a term may be applied to each. In some cases the strength of mechanism may be expressed quantitatively as a natural consequence of the technology (e.g., the number of bits in a CRC, the particular function employed); this quantitative measure of strength may be employed as the basis for rating.

The Part II evaluations may also be expected to exhibit a greater sensitivity to technological advances than the Part I evaluations. This sensitivity derives from the present empirical basis of some of the Part II security services as compared to the theoretical foundation of Part I. Research advances may help change this situation. As the state-of-the-art advances, the threshold for high evaluations may also be expected to increase. Therefore, a rating may become dated and may change upon reevaluation.

In general, mechanisms that only protect against accidents and malfunctions cannot achieve an evaluation of strength of mechanism above minimum. Mechanisms must provide protection against deliberate attacks in order to obtain at least a good evaluation.

The summary report of a network product will contain the rating reflecting the Part I evaluation plus a paired list of Part II services and the evaluation for each. For example, network product XYZ might be rated as follows: [B2, security service-1: minimum, security service-2: not offered, security service-3: none, ... ,security service-n: (functionality: good, strength of mechanism: fair, assurance: good)]. In some cases where the security service is addressed outside this document (e.g., COMSEC), the evaluation from the external source may be reflected in the evaluation report. In such cases, the terms used will differ from those listed above.

5.4. Relationship to ISO OSI-Architecture

An effort is underway to extend the ISO Open System Interconnection (OSI) architecture by defining "general security-related architectural elements which can be applied appropriately in the circumstances for which protection of communications between open systems is required." - Familiarity with OSI terminology is assumed in this discussion. The scope of this security addendum "provides a general description of security services and related mechanisms,

- ISO 7498/Part 2 - Security Architecture, ISO / TC

97 / SC 21 / N1528 / WG 1 Ad hoc group on Security, Project 97.21.18, September 1986.

which may be provided by the Reference Model; and defines the positions within the Reference Model where the services and mechanisms may be provided."

There is considerable overlap between the OSI Security Addendum and Part II. At the time of writing, the OSI document is evolving, making it difficult to exactly define the relationship. Therefore, the following statements may have to be modified in the future.

Some of the security services identified in the OSI Security Addendum are covered by Part I of this Interpretation; others are addressed in Part II. The emphasis is on making sure that all services are covered. The distinction between the security service and the mechanism that implements the service is less strong in this Interpretation than

in the OSI Security Addendum. The OSI Addendum generally addresses Functionality, occasionally addresses Strength of Mechanism, and rarely addresses Assurance, while in this Interpretation, especially in Part I, assurance is a major factor.

The scope of the OSI Security Addendum is limited: "OSI Security is not concerned with security measures needed in end systems, installations and organizations except where these have implications on the choice and position of security services visible in OSI." The TCSEC and this Interpretation include OSI concerns as a proper subset.

5.5. Selecting Security Services for a Specific Environment

The enumeration of security services in Part II is representative of those services that an organization may choose to employ in a specific network for a specific environment. But not all security services will be equally important in a specific environment, nor will their relative importance be the same among different environments. The network management has to decide whether the rating achieved by a network product for a specific criterion is satisfactory for the application environment.

As an abstract example, consider the network product XYZ which has received the rating [B2, security service-1: minimum, security service-2: not offered, ...]. The management of network K may decide that they do not require security service-2, so the absence of this service does not effect the acceptability of the XYZ product; however, the management of network Q may decide that security service-2 is essential, so the absence of this service disqualifies product XYZ. The management of network P may decide security service-1 is very important and that any rating less than good is unacceptable, thereby disqualifying product XYZ; while the management of network R may decide that security service-1 need only be rated minimum.

As a more concrete example, consider an application environment where wire-tapping is not a threat, such as aboard an airplane or in an underground bunker. A Local Area Network (LAN) in such an environment can be physically protected to the system-high security mode without encryption because the system exists within a protected perimeter. In such environments, management may decide that labeling and access control based on labels provide sufficient protection if sufficient mechanisms exist to protect the integrity of the labels. Cryptographic mechanisms are deemed unnecessary. By way of contrast, when the LAN environment involves passage through unprotected space, management may decide that a LAN must provide integrity protection employing a cryptographic mechanism.

6. General Assurance Approaches

This section addresses assurance approaches applicable to many security services.

The logic of the protocols and the implementation of countermeasures may be shown correct and effective by formal methods where possible (i.e., where tools exist) and informal ones otherwise.

To provide assurance that the security service can respond to various forms of external attacks, various methods of real and simulated testing can be applied, including:

1. Functional testing
2. Periodic testing
3. Penetration testing
4. Stress testing
5. Protocol testing for deadlock, liveness, and other security properties of the protocol suites

In addition, the trusted computer base provides an execution environment that is extremely valuable in enhancing the assurance of a variety of security services. The discretionary and mandatory access controls can be employed in the design and implementation of these services to segregate unrelated services. Thus, service implementation that is complex and error-prone or obtained from an unevaluated supplier can be prevented from degrading the assurance of other services implemented in the same component. Furthermore, a TCB ensures that the basic protection of the security and integrity- of the information entrusted to the network is not diluted by various supporting security services identified in this Part II. See also the discussion of Integrity in the Supportive Primitives section.

In general, assurance may be provided by implementing these features in a limited set of subjects in each applicable NTCB partition whose code and data have a unique mandatory integrity level to protect against circumvention and

tampering.

- See, for example, Biba, K.J., "Integrity Consideration for Secure Computer Systems," ESD-TR-76-372, MTR-3153, The MITRE Corporation, Bedford, MA, April 1977.

Assurance of trustworthiness of the design and implementation of Part II mechanisms may be related to the assurance requirements in Part I. The following factors are identified as contributing to an assurance evaluation: service design and implementation, service testing, design specification and verification, configuration management, and distribution.

6.1. Service Design and Implementation Factors

An evaluation rating of fair indicates that the implementation of the service employs the provisions of the TCB for a distinct address space. In addition, the implementation of the service is internally structured into well-defined largely independent modules; makes effective use of available hardware to separate those elements that are protection-critical to the service from those that are not; is designed such that the principle of least privilege is enforced; and the user interface is completely defined and all elements relevant to the service are identified.

An evaluation rating of good indicates that the service, in addition, incorporates significant use of layering, abstraction and data hiding; and employs significant system engineering directed toward minimizing complexity and separating modules that are critical to the service.

6.2. Service Testing Factors

With respect to security testing, an evaluation of minimum indicates that the service was tested and found to work as claimed in the system documentation; that testing was done to assure that there are no obvious ways for an unauthorized user to bypass or otherwise defeat the security constraints and objectives; and that testing included a search for obvious flaws that would allow inconsistent or improper modification of data used by the service, either by external software or by errors in the implementation of the service.

An evaluation rating of fair indicates that, in addition to the minimum factors, a team of individuals who thoroughly understand the specific implementation subjected its design documentation, source code, and object code to through analysis and testing with the objectives of uncovering all design and implementation flaws that would permit a

subject external to the NTCB to defeat the purposes of the service. A fair system is relatively resistant to defeat of the purpose of the service. A fair evaluation indicates that all discovered flaws were removed or neutralized and the system retested to demonstrate that they have been eliminated and that new flaws have not been introduced. Testing demonstrates that the service implementation is consistent with the specification.

An evaluation rating of good indicates that, in addition to the fair factors, the system is more resistant to defeat of service; and that no design flaws and no more than a few correctable implementation flaws were found during testing and there is reasonable confidence that few remain. Manual or other mapping of the specifications to the source code may form a basis for testing.

6.3. Design Specification and Verification Factors

With respect to design specification and verification, an evaluation rating of minimum indicates that an informal model of the properties of the service is maintained over the life cycle of the system. Additional requirements for an evaluation rating of fair have not been defined.

An evaluation rating of good indicates that, in addition, a formal model of the properties of the service is maintained over the life cycle of the system and demonstrated to be consistent with its axioms; and that a descriptive specification of the service-relevant code is maintained that completely and accurately describes it in terms of exceptions, error messages, and effects.

6.4. Configuration Management Factors

With respect to configuration management, an evaluation rating of minimum indicates that during development and maintenance of the service, a configuration management system was in place that maintained control of changes to specifications, other design data, implementation documentation, source code, the running version of the object code, test fixtures, test code, and documentation.

An evaluation rating of fair indicates that, in addition, the configuration management system assures a consistent mapping among all documentation and code associated with the current version of the system; and for comparing a newly generated version with the previous version in order to ascertain that only the intended changes have been made in the code.

An evaluation rating of good indicates that, in addi-

tion, configuration management covers the entire life-cycle; that it applies to all firmware, and hardware that supports the service; and that a combinations of technical, physical, and procedural safeguards are used to protect from unauthorized modification or destruction the master copy or copies of all material used to generate the implementation of the service.

6.5. Distribution Factors

There are currently no requirements for minimum and fair evaluation ratings.

With respect to distribution, an evaluation rating of good indicates that a control and distribution facility is provided for maintaining the integrity of the mapping between the master data describing the current version of the service and the master copy of the code for the current version. Procedures (e.g., site security acceptance testing) shall exist for assuring that the software, firmware, and hardware updates distributed are exactly as specified by the master copies.

7. Supportive Primitives

This subsection describes mechanisms and assurance techniques that apply across multiple security services. They are grouped together here for convenience and are referenced in the appropriate service subsections of Section 8.

Encryption is a pervasive mechanism for many security services; protocols are of fundamental essence in networks. The information in this Section 7 is provided as background and support for the services addressed in Section 8.

7.1. The Encryption Mechanism

7.1.1. Functionality Factors

Encryption is a tool for protecting data from compromise or modification attacks. Through its use, release of message content and traffic analysis can be prevented; message stream modification, some denial of message service, and masquerading can be detected. For example, an ISO document-, describing the use of encipherment techniques in communications architectures, has been published as a U.S. member body contribution for consideration as cryptographic security protection in the Open System Interconnection environment. Encryption is probably the most important and

widely used security mechanism when there is a wiretap threat; sometimes it is even confused with being a service.

Use of the encryption mechanism leads to a requirement for key management (e.g., manually or in the form of key distribution protocols and key-distribution centers.)

7.1.2. Strength of Mechanism Factors

The strength of a cryptographic cipher is determined by mathematical and statistical analysis; the results are typically expressed in the workfunction required for unauthorized decryption. In many cases this analysis is classified; the results are available only as a statement of the highest level of classified data which may be protected by use of the mechanism.

When encryption is used in networks, it may be combined with network protocols to protect against disclosure. The strength of the ciphers, the correctness of the protocol logic, and the adequacy of implementation, are primary factors in assessing the strength of Data Confidentiality using

- Addendum to the Transport Layer Protocol Definition
for Providing Connection Oriented End-to-End Crypto-
graphic Data Protection Using A 64-bit Block Cipher,
ISO TC 97 / SC 20 / WG 3, N 37, January 10, 1986.

cryptography techniques. Algorithms are characterized by the National Security Agency on a pass/fail basis in terms of the sensitivity of the information which the encryption algorithm is approved to protect.

7.1.3. Assurance Factors

The analysis of encryption techniques is quite different from the formal specification and verification technology employed as the basis of trust in the TCSEC. Much of this analysis is classified. Consequently, assurance of encryption techniques will be provided by the National Security Agency. Normally, a separate assurance rating will not be given.

7.2. Protocols

7.2.1. Functionality Factors

Protocols are a set of rules and formats (semantic and syntactic) that determine the communication behavior between entities in a network. Their design and implementation is crucial to the correct, efficient, and effective transfer of information among network systems and sub-systems.

Many network security services are implemented with the help of protocols, and failures and deficiencies in the protocol result in failures and deficiencies in the security service supported by the protocol.

One class of design, or logical, deficiencies in protocols are those having some form of denial of service as a consequence. This class includes deadlocks, livelocks, unspecified receptions, lack-of-liveness, and non-executable interactions. A protocol with one of these design flaws can cease to function under circumstances that can occur during normal operation but which were not anticipated by the designer. Such flaws result in trapping the protocol into states that are nonproductive or which cause the protocol to halt or have unpredictable effects.

Another class of design concerns are typical of protocols that must work despite various kinds of random interference or communication difficulties, such as noise, message loss, and message reordering. It should be noted that most networks are designed in a layered fashion, in which each protocol-based service is implemented by invoking services available from the next lower layer. This means that if one layer provides protection from certain types of communication difficulties, higher layers need not address those problems in their design.

A third class of design deficiency might occur in protocols that are expected to work in the presence of malicious interference, such as active wiretapping. Such protocols should have countermeasures against Message Stream Modification (MSM) attacks.

7.2.2. Strength of Mechanism Factors

Protocol deficiencies may lie either in their design or their implementation. By an implementation deficiency is meant a lack of correspondence between a protocol specification and its implementation in software.

7.2.3. Assurance Factors

Assurances of implementation correctness may be addressed by techniques such as design specification and verification, and testing.

Ideally, all of the network protocol functions would be verified to operate correctly. However, verification of large amounts of code is prohibitively expensive (if not impossible) at the current state-of-the-art, so the code to be verified must be kept to an absolute minimum. It seems feasible to split up a complex protocol (e.g., the TCP) into trusted portions (i.e., the software that performs security-related functions) and untrusted portions (i.e., other software) so that only the security-related portions must be shown to meet the requirements of Part I. However, there is a general concern about the extent to which trusted portions of protocols can be identified and protected from untrusted portions.

Methods for assuring the design correctness of protocols involve the use of tools and techniques specially oriented toward the kinds of problems peculiar to protocols. Either formal methods, or testing, or both, may be used.

Some assurance in design correctness may be obtained simply by basing the protocol design on a well-understood model or technique found in the literature if it is known to address the kinds of problems likely to arise. This assurance is lessened to the extent that the actual protocol differs from the published version.

7.2.3.1. Formal Methods

Formal techniques of protocol definition and validation have advanced to the point that they may be applied to actual protocols to verify the absence of deadlocks, livelocks, and incompleteness for design verification. When the state-of-the-art of formal tools is inadequate, or when the sponsor decides not to employ formal tools, informal methods may be used. The evaluation of protocol specification and verification should indicate which assurance tools have been employed.

Formal methods for protocol specification and verification are typically based on a finite-state machine concept, extended in one of various ways to represent the concurrency and communication properties characteristic of networks. Communicating sequential machines and Petri nets have been used as a functional modeling context for protocols, and experimental automated verification tools based on these models have been developed. Different models and tools may need to be used depending on the design objective for which assurance is desired.

To the extent that the protocol model and implementation permit separation by layers, the functional model, proofs, demonstration, and arguments may optionally be applied to individual layers or sets of adjacent layers. Generally, the assurances obtained about protocols in one layer are conditional on, or relative to, assurances for protocols in lower layers.

7.2.3.2. Testing

Protocol testing is another method to assure the correctness of the protocols other than formal verification. Protocol testing has been employed to certify implementation conformance to standards such as X.25, TCP, and TP4 with a moderate level of success.

The type of testing called for can be referred to as conformance testing and penetration testing. The purpose of performing these tests is to obtain a moderate level of confidence on the correct operation of the protocols.

Objectives should be to uncover design and implementation flaws that would cause the protocols to perform their functions incorrectly, and to determine if the Message Stream Modification (MSM) countermeasures are effective, if applicable. They may attempt to uncover all kinds of logical deficiencies, such as deadlocks, livelocks, unspecified receptions, lack-of-liveness, and non-executable interactions. All discovered flaws should be corrected and the implementations retested to demonstrate that they have been eliminated and that new flaws have not been introduced. For a successful conclusion to a test suite, no design flaws and no more than a few correctable implementation flaws may be found during testing, and there should be reasonable confidence that few remain. Manual or other mapping of the protocol specification to the source code may form a basis for testing.

Protocols should be thoroughly analyzed and tested for their responses to both normal and abnormal data type messages. Testing must be done for both normal and degraded mode of operation both in controlled environment and in the environment of deployment.

8. Documentation

The section headings in these Part II Documentation criteria are the same as those employed for Part I Documentation criteria. The documentation produced in response to both sets of criteria may optionally be combined or published separately, as the sponsor sees fit.

8.1. Security Features User's Guide

A single summary, chapter, or manual in user documentation shall describe the Part II security services, guidelines on their use, and how they interact with one another.

This user documentation describes security services at the global (network system) level, at the user interface of each component, and the interaction among these.

8.2. Trusted Facility Manual

A manual addressed to the network and component subsystem administrator shall present cautions about functions and privileges that should be controlled to maintain network security. The manual shall describe the operator and administrator functions related to security services. It shall provide guidelines on the consistent and effective use of the network security services, how they interact, and facility procedures, warnings, and privileges that need to be controlled in order to maintain network security.

The software modules that provide security services shall be identified. The procedures for secure generation of new security service object modules from source after modification of source code shall be described. It shall include the procedures, if any, required to ensure that the network is initially started in a secure manner. Procedures shall also be included to resume secure system operation after any lapse in operation.

This manual shall contain specifications and procedures to assist the system administrator to maintain cognizance of the network configuration. These specifications and procedures shall address:

1. The implications of attaching new components to the network.
2. The case where certain components may periodically leave the network (e.g., by crashing or by being disconnected) and then rejoin.
3. Incremental updates; that is, it must explicitly indicate which security services may change without others also changing.

The physical and administrative environmental controls shall be specified. Any assumptions about security of a given network should be clearly stated (e.g., the fact that all communications links must be physically protected to a certain level).

8.3. Test Documentation

A document shall be provided that describes the test plan and test procedures that show how the security services were tested, and results of the security services' functional testing.

The description of the test plan should establish the context in which the testing was or should be conducted. The description should identify any additional test components that are not part of the system being evaluated. This includes a description of the test-relevant functions of such test components, and a description of the interfacing of those test components to the system being evaluated. The description of the test plan should also demonstrate that the tests adequately cover the network security policy. The tests should also include network configuration and sizing.

As identified in Appendix A, the entity being evaluated may be a networking subsystem to which other components must be added to make a complete network system. In that case, test documentation must include contextual definition because, at evaluation time, it is not possible to validate the test plans without the description of the context for testing the networking subsystem.

8.4. Design Documentation

Documentation shall be available that provides a description of the network philosophy of protection and an explanation of how this philosophy is translated into the security services offered. The interfaces between security services shall be described. The security policy also shall be stated.

The system description addresses the network security architecture and design by specifying the security services in the network, and in what way they must cooperate to support network security objectives. If a network supports a set of security policies and permits components with different policies to communicate, the relationships between the policies shall be defined.

9. Specific Security Services

This section contains specific security services that may be provided in networks. The structure of the specific security services in the balance of Part II is represented in Table II-2. This table shows the network security con-

cerns addressed, the criteria for each concern, and the evaluation range for each criterion.

9.1. Communications Integrity

Communications integrity is a collective term for a number of security services. These services, described below, are all concerned with the accuracy, faithfulness, non-corruptibility, and believability of information transfer between peer entities through the computer communications network.

Integrity is an important issue. However, there is considerable confusion and inconsistency in the use of the term. The term is used to address matters such as consistency, accuracy, concurrency, and data recovery, modification access control (write, append, delete, update) and the credibility of information that is read by a process.-

The mechanisms that can be used to enforce communication integrity have some strong similarities to the mechanisms that are used to enforce discretionary and mandatory access controls. Integrity in Part I is concerned with access control, specifically the ability of subjects to modify objects. This should be contrasted with the Part II concerns for communications integrity described below.

9.1.1. Authentication

+ Functionality

The network should ensure that a data exchange is established with the addressed peer entity (and not with an entity attempting a masquerade or a replay of a previous establishment). The network should assure that the data source is the one claimed. When this service is provided in support of a connection-oriented association, it is known as Peer Entity Authentication; when it supports a connectionless association, it is known as Data Origin Authentication.

Attempts to create a session under a false identity or playing back a previous legitimate session initiation sequence are typical threats for which peer entity

- See, for example, Biba, K.J., "Integrity Consideration for Secure Systems," ESD-TR-76-372, MTR-3153, The Mitre Corporation, Bedford, MA, April 1977; and Juene-man, R. R., "Electronic Document Authentication," IEEE

Network Magazine, April 1987, pp 17-23.

authentication is an appropriate countermeasure.

Authentication generally follows identification, establishing the validity of the claimed identity providing protection against fraudulent transactions. Identification, authentication, and authorization information (e.g., passwords) should be protected by the network.

Available techniques which may be applied to peer authentication mechanisms are:

1. Something known by the entity (e.g., passwords)
2. Cryptographic means
3. Use of the characteristics and/or possessions of the entity

The above mechanisms may be incorporated into the (N)-layer peer-to-peer protocol to provide peer entity authentication.

To tie data to a specific origin, implicit or explicit identification information must be derived and associated with data. Ad hoc methods exist for authentication which may include verification through an alternate communications channel, or a user-unique cryptographic authentication.

When encryption is used for authentication service, it can be provided by encipherment or signature mechanisms. In conventional private-key cryptosystems, the encryption of a message with a secret key automatically implies data origin authenticity, because only the holder of that key can produce an encrypted form of a message. However, the kind of authentication provided by the conventional private-key cryptosystem can protect both sender and receiver against third party enemies, but it cannot protect against fraud committed by the other. The reason is that the receiver knowing the encryption key, could generate the encrypted form of a message and forge messages appearing to come from the sender. In the case where possible disputes that may arise from the dishonesty of either sender or receiver, a digital signature scheme is required.

In public-key cryptosystems, message secrecy and message/sender authenticity are functionally independent. To achieve authenticity, the message is "decrypted" with the secret key of the sender to provide proof of its origin, but that does not conceal the message. If both secrecy and authenticity are required, a public-key signature scheme must be used. This subject is extensively addressed in the ISO/CCITT context of Directory authentication=.

= The Directory - Authentication Framework (Melbourne, April 1986), ISO/CCITT Directory Convergence Document #3.

Basis for Rating: Presence or absence.

Evaluation Range: None or present.

+ Strength of Mechanism

The security provided by the passwords mechanism is very sensitive to how passwords are selected and protected. The security provided by a password depends on composition, lifetime, length, and protection from disclosure and substitution. Password Management Guidance is contained in a separate document-.

When cryptographic techniques are used, they may be combined with "hand-shaking" protocols and "liveness" assurance procedures to protect against masquerading and replay. The "liveness" assurance procedures may be provided by:

1. Synchronized clocks
2. Two and three ways handshakes
3. Non-repudiation services provided by digital signature and/or notarization mechanisms

The strength of the ciphers, the correctness of the protocol logic, and the adequacy of implementation are three primary factors in assessing the strength of authentication using cryptography techniques. See also the Encryption Mechanism section.

Basis for Rating: In order to obtain a rating of good using passwords, such usage must conform to Password Management Guidance-. The strength of a cryptographic mechanism will be provided by the National Security Agency.

Evaluation Range: None to good.

+ Assurance

Basis for rating assurance is concerned with guaranteeing or providing confidence that features to address authen-

tication threats have been implemented correctly and that the control objectives of each feature have been actually and accurately achieved.

This assurance may be addressed by analysis of the strength of the authentication exchange mechanism. This includes password scheme and/or cryptographic algorithm analysis and the automated protocol testing for deadlock,

- Department of Defense Password Management Guide-
line, National Computer Security Center, CSC-STD-002-
85, 12 April 1985.

liveness, and other security properties of the "hand-shaking" protocols.

Many of the assurance approaches are common to other security services. See the General Assurance Approaches section for further information. Cryptographic mechanisms may be employed for peer entity authentication. These mechanisms, and their assurance, are discussed in a separate section.

Basis for Rating: See the General Assurance Approaches section.

Evaluation Range: None to good.

9.1.2. Communications Field Integrity

Communications Field Integrity refers to protection of any of the fields involved in communications from unauthorized modification. Two well-known fields are the protocol-information (a.k.a. header) field and the user-data field. A protocol-data-unit (PDU) (a.k.a. packet, datagram) always contains protocol-information; user-data is optional.

Other division and identification of fields are possible. Some communications systems identify such fields as control and priority. For generality, this section refers to any field as containing data; this data may in fact be protocol-information or the contents of some other identified field. For convenience, the term data integrity will be used synonymously with communications field integrity. Data integrity may be provided on a selective field basis; some selection may be made by the network architect, some selection may be made by the network administration, and some may be left to the user.

It should be mentioned that in a layered protocol the combination of layer N protocol-information plus layer N user-data is considered to be all user-data in layer N-1. It is also important to note the definition of a message and the relationship between PDUs and messages. Each PDU may constitute an independent message, or a sequence of PDUs may constitute a single message.

+ Functionality

Data integrity service counters active threats and protects data against unauthorized alteration. The network should ensure that information is accurately transmitted from source to destination (regardless of the number of intermittent connecting points). The network should be able to counter both equipment failure as well as actions by persons and processes not authorized to alter the data. Protocols that perform code or format conversion will preserve the integrity of data and control information.

The network should also have an automated capability of testing for, detecting, and reporting errors that exceed a threshold.

Since communication may be subject to jamming/spoofing attack, line and node outages, hardware and software failures, and active wiretapping attacks, there should exist effective countermeasures to counter possible communications threats. The countermeasures may include policy, procedures, automated or physical controls, mechanisms, and protocols means.

Basis for Rating: Data Integrity service may be evaluated according to its ability to detect integrity violations. The following progression relates features and evaluation.

Functionality would be evaluated as minimum if either of the following two levels of features were provided:

1. Integrity of a single connectionless PDU. This takes the form of determining whether a received PDU has been modified.
2. Integrity of selected fields within a connectionless PDU. This takes the form of determining whether the selected fields have been modified.

Functionality would be evaluated as fair if, in addition, either of the following two levels of features were provided:

1. Integrity of selected fields transferred over a connection. This takes the form of determining whether the selected fields have been modified, inserted, deleted, or replayed.
2. Integrity of all user-data on a protocol layer connection. This service detects any modification, insertion, deletion, or replay of any PDU of an entire PDU sequence with no recovery attempted.

Functionality would be evaluated as good if, in addition, the following feature is provided:

Integrity of all user-data on a protocol layer connection. This service detects any modification, insertion, deletion, or replay of any PDU of an entire PDU sequence with recovery attempted.

Evaluation Range: None to good.

+ Strength of Mechanism

Policy, procedures, automated or physical controls, mechanisms, and protocols should exist for ensuring that data has not been subject to excessive random errors and unauthorized message stream modification, such as alteration, substitution, reordering, replay, or insertion. Message stream modification (MSM) countermeasures should be identified and shown to be effective. A technology of adequate strength should be selected to resist MSM.

The probability of an undetected error should be specified as an indication of the strength of mechanism. The network should have an automated capability for testing, detecting, reporting, and/or recovering from those communication errors/corruptions that exceed specified network requirements.

Basis for Rating: When encryption is used in networks, it is combined with network protocols to protect against unauthorized data modification. The strength of the ciphers, the correctness of the protocol logic, and the adequacy of implementation are three primary factors in assessing the strength of Data Integrity using cryptography techniques. See the Encryption Mechanism section for further information.

Evaluation Range: None to good.

+ Assurance

Basis for rating: Assurance is concerned with guaranteeing or providing confidence that features to address Data Integrity threats have been implemented correctly and that the control objectives of each feature have been actually and accurately achieved.

Many of the assurance approaches for data integrity are common to other security services. See the General Assurance section for further information.

Evaluation Range: None to good.

9.1.3. Non-repudiation

+ Functionality

Non-repudiation service provides unforgeable proof of shipment and/or receipt of data.

This service prevents the sender from disavowing a legitimate message or the recipient from denying receipt. The network may provide either or both of the following two forms:

1. The recipient of data is provided with proof of origin of data that will protect against any attempt by the sender to falsely deny sending the data or its contents.
2. The sender is provided with proof of delivery of data such that the recipient cannot later deny receiving the data or its contents.

Basis for Rating: Presence or absence of each of the two forms.

Evaluation Range: None or present.

Discussion: Digital signatures are available techniques that may be applied to non-repudiation mechanisms. Digital signature mechanisms define two procedures:

1. Signing a data unit
2. Verifying a signed data unit

The signing process typically employs either an encipherment of the data unit or the production of a cryptographic checkfunction of the data unit, using the signer's

private information as a private key.

The verification process involves using the public procedure and information to determine whether the signature was produced with the signer's private key.

It is essential that the signature mechanism be unforgeable and adjudicable. This means that the signature can only be produced using the signer's private information, and in case the signer should disavow signing the message, it must be possible for a judge or arbitrator to resolve a dispute arising between the signer and the recipient of the message.

Digital signature schemes are usually classified into one of two categories: true signatures or arbitrated signatures. In a true signature scheme, signed messages produced by the sender are transmitted directly to the receiver who verifies their validity and authenticity. In an arbitrated signature scheme, all signed messages are transmitted from the sender to the receiver via an arbitrator who serves as a notary public. In the latter case, a notarization mechanism is needed.

Both public-key and conventional private-key cryptosystems can be utilized to generate digital signatures. When a message M is to be signed by a private-key cryptosystem, the signature is a computed quantity catenated to M and sent along with it. In a public-key implementation, when a message M is to be signed, a transformation using the secret key is applied to M before transmitting it. Thus, the signature is presented by the resulting transformed message.

+ Strength of Mechanism

Basis for Rating: The strength and trustworthiness given to non-repudiation service is bounded by the trust in the underlying cryptography implementing digital signature mechanism, the correctness of the protocol logic, and the adequacy of protocol implementation. Additional information may be found in the separate sections addressing these subjects.

Evaluation Range: None to good.

+ Assurance

Basis for Rating: Assurance is concerned with guaranteeing or providing confidence that features to provide non-repudiation service have been implemented correctly and that the control objectives of each feature have been actu-

ally and accurately achieved.

This assurance is addressed by analysis of the logic of the protocols and the implementation of the digital signature mechanisms to show correctness and effectiveness by formal methods where possible (i.e., where tools exist) and informal ones otherwise.

The information in the General Assurance, Encryption Mechanisms, and Protocols sections also applies.

Evaluation Range: None to good.

9.2. Denial of Service

Assurance of communications availability would probably be more properly identified as a service, while denial-of-service (DOS) would be identified as a threat. However, it is traditional to employ denial of service as the identifier of this topic.

DOS detection is highly dependent on data integrity checking/detection mechanisms. Other mechanisms relating to data ordering, modification, loss, or replay (e.g., sequence numbers, frame counts) are also measures of DOS protection.

A denial-of-service condition exists whenever the throughput falls below a pre-established threshold, or access to a remote entity is unavailable. DOS also exists when resources are not available to users on an equitable basis. Priority and similar mechanisms should be taken into account in determining equity. If a connection is active, a DOS condition can be detected by the maximum waiting time (MWT) or the predetermined minimum throughput. However, when a connection is quiescent, a protocol entity at one end of a connection has no way of determining when the next packet should arrive from its corresponding peer entity. It is thus unable to detect a DOS attack that completely cuts off the flow of packets from that entity.

Denial of service conditions should be considered for all services being provided by the network. As discussed below for specific services, depending on the strength of mechanism the network should be able to detect, recover, and/or resist denial of service conditions. The specific conditions, which the network will address, are determined through the use of informal models, such as Mission(s) Model, Threat Model, Life Cycle Model, and Service Oriented Model. The network manager or sponsor shall determine the network's denial of service requirements and shall establish the desired service criteria accordingly.

9.2.1. Continuity of Operations

+ Functionality

The security features providing resistance against DOS external attacks and the objectives that each feature will achieve may include the following:

1. Use of active or passive replacement or other forms of redundancy throughout the network components (i.e., network nodes, connectivity, and control capability) may enhance reliability, reduce single-point-of-failure, enhance survivability, and provide excess capacity.
2. Reconfiguration to provide network software maintenance and program down-loading to network nodes for software distribution, and to provide initialization and reconfiguration after removing failed or faulty components and replacing with replaced components can isolate and/or confine network failures, accommodate the addition and deletion of network components, and circumvent a detected fault.
3. Distribution and flexibility of network control functions utilizing a distributed control capability to reduce or eliminate the possibility of disabling the network by destroying or disabling one or a few network control facilities, and a flexible control capability which is able to respond promptly to emergency needs, such as increase in traffic or quick restoration, can improve the capability to respond promptly to the changes in network topology and network throughput thereby enhancing survivability and continuity of operation, perhaps by enforcing precedence and preemption on traffic handling.
4. Fault tolerance mechanisms provide a capability to deal with network failures and to maintain continuity of operations of a network including the following features: error/fault detection, fault treatment, damage assessment (analysis on effects of failures), error/failure recovery, component/segment crash recovery, and whole network crash recovery.
5. Security controls could include community of interest separation through creation of logical subnets with disjoint non-hierarchical mandatory access control categories, and protection of control information from active wiretapping.

Basis for Rating: The network should ensure some minimum specified continuing level of service. The follow-

ing service would be considered minimum:

- a) Detect conditions that would degrade service below a pre-specified minimum and would report such degradation to its operators.

The following service would be considered fair:

- b) Service that would continue in the event of equipment failure and actions by persons and processes not authorized to alter the data. The resiliency may be provided by redundancy, alternate facilities, or other means. The service provided may be degraded and/or may invoke priorities of service.

The following service would be considered good:

- c) The same as (a), but with automatic adaptation.

Evaluation Range: None to good.

+ Strength of Mechanism

Network operational maintenance is based on mechanisms whose robustness may decrease inversely with network loading. It may be nearly impossible to guarantee sufficiently robust testing, regardless of whether done off-line with simulated loading or operationally.

In addition to rigorous analysis to assure algorithmic correctness in dealing with the "internal failures" (e.g., component, segment, or system failures caused by errors in resource allocation policy or mechanism implementation), countermeasures shall also be employed against "external attacks" such as physical attacks.

Basis for Rating: For each DOS feature defined above, it is possible to assign a rating such as none, minimum, fair, and good for the assessment of a network's "DOS strength" with respect to that particular feature.

For example, major ways of providing fault-tolerant mechanisms include:

1. Error/fault detection
2. Fault treatment
3. Damage assessment (analysis on effects of failures)
4. Error/failure recovery

5. Component/segment crash recovery
6. Whole network crash recovery

Evaluation Range: None to good.

+ Assurance

Assurance is concerned with guaranteeing or providing confidence that features to address DOS threats have been implemented correctly and that the objectives of each feature have been actually and accurately achieved.

This assurance may be addressed by analysis for weakness or anomalous behavior of the resource allocation policy/mechanisms of the network using various formal models such as queuing theoretic models, hierarchical service models, protocol models, or resource allocation models which can be analyzed for deadlock, liveness, and other security properties.

Basis for Rating: To provide assurance that the network can respond to various forms of denial of service conditions, the following methods may be employed:

1. Simulation
2. Testing
 - a. Functional
 - b. Periodic
 - c. Penetration
3. Measurement under extreme conditions

Distribution, as discussed as one of the General Assurance Factors, can increase the assurance that the software deployed is authentic and appropriate in the context of deployment of new software and in crash recovery. In addition, development in a closed environment can increase assurance.

Evaluation Range: None to good.

9.2.2. Protocol Based DOS Protection Mechanisms

+ Functionality

Mechanisms for addressing DOS are often protocol based

and may involve testing or probing. Any communications availability service should consider using existing communications protocol mechanisms where feasible so as not to increase network overhead. DOS mechanisms add overhead that may have some adverse impact on network performance. The benefits of value-added functions should offset the resultant performance cost.

For example, in order to detect throughput denial of service, a process may exist to measure the transmission rate between peer entities under conditions of input queuing. The measured transmission rate shall be compared with a predetermined minimum to detect a DOS condition and activate an alarm.

Another example is a protocol to detect failure to respond within a predetermined time between peer entities. This protocol would determine the remote entity's ability to respond to the protocol.

A request-response mechanism such as "are-you-there" message exchange may be employed to detect DOS conditions when the connection is quiescent. The request-response mechanism involves the periodic exchange of "hello", and "are-you-there" messages between peer entities to verify that an open path exists between them; such a mechanism should be protected against selective message passing. Based on the ability to respond and the response time to the request-response mechanism, the "availability" of a remote entity can be determined and the DOS condition can be detected.

Request-response mechanisms have been known to crash networks when coupled with hardware failures and/or abnormal loading. Incompatibilities also sometimes show up when dissimilar networks are interconnected. Any polling sequence should probably be metered to prevent creating a DOS condition.

Basis for Rating: The number of protocol based mechanisms could be used for evaluation. If only one mechanism were provided, the functionality might be rated as minimum. If two or three mechanisms were provided, the functionality might be rated as fair. If more than three mechanisms were provided, the functionality might be rated as good.

Evaluation Range: None to good.

+ Strength of Mechanism

Basis for Rating: Network protocol robustness may decrease inversely with network loading. Testing, off-line with simulated loading or operationally, and rigorous

analysis to assure protocol correctness in dealing with the "internal failures" and against "external attacks" are appropriate ways of establishing strength of mechanism.

Evaluation Range: None to good.

+ Assurance

Assurance is concerned with guaranteeing or providing confidence that features to address DOS threats have been implemented correctly and that the objectives of each feature have been actually and accurately achieved.

Basis for Rating: This assurance may be addressed by analysis for weakness or anomalous behavior of the network protocols using various formal models such as queuing theoretic models, hierarchical service models, petri nets, or resource allocation models which can be analyzed for deadlock, liveness, and other security properties.

To provide assurance that the network can response to various forms of external attacks, the following methods may be employed:

1. simulation
2. testing
 - functional
 - periodic
 - penetration
3. measurement under extreme conditions

Distribution, as discussed as one of the General Assurance Factors, can increase the assurance that the software deployed is authentic and appropriate in the context of deployment of new software and in crash recovery. In addition, development in a closed environment can increase assurance.

Evaluation Range: None to good.

9.2.3. Network Management

+ Functionality

DOS resistance based on a system/message integrity

measure is two-tiered. Tier one deals with communications protocols. Tier two addresses network management (and maintenance). These tiers for the most part operate independently.

Network management and maintenance in tier two deals with network health, detecting failures and overt acts that result in denial or reduced service. Simple throughput may not necessarily be a good measure of proper performance. Loading above capacity, flooding, replays, and protocol retry due to noise in the channel can reduce service below an acceptable level and/or cause selective outages. Management protocols, such as those which configure the network or monitor its performance, are not described well by the existing protocol reference models.

A DOS attack may cause disruption of more than one peer entity association. For this reason detection and correction may be implemented in tier two. The detection of a potential DOS condition by a peer entity should be reported by the layer management functions of those entities. The determination of a DOS attack is an application management function, and the corrective action is a system management function.

Basis for Rating: Presence or absence.

Evaluation Range: None or present.

+ Strength of Mechanism

Network operational maintenance is based on mechanisms whose robustness may decrease inversely with network loading (e.g., update of routing tables).

Basis for Rating: Integrity and adequacy of control in a network are the keys in coping with denial of service conditions. In addition to rigorous analysis to assure algorithmic correctness in dealing with the "internal failures" (e.g., component, segment, or system failures caused by errors in resource allocation policy or mechanism implementation), countermeasures shall also be employed against "external attacks," such as physical attacks and attacks against network control.

Based on these characterizations, a set of ratings can be assigned to each category under the fault tolerance feature and an overall rating can then be determined for a network's strength in providing "fault tolerance mechanisms".

Evaluation Range: None to good.

+ Assurance

Basis for Rating: Assurance may be addressed by analysis for weakness or anomalous behavior of the network management policy/mechanisms of the network using various formal models such as queuing theoretic models, hierarchical service models, protocol models, or resource allocation models which can be analyzed for deadlock, liveness, and other security properties.

Distribution, as discussed as one of the General Assurance Factors, can increase the assurance that the software deployed is authentic and appropriate in the context of deployment of new software and in crash recovery. In addition, development in a closed environment can increase assurance.

Evaluation Range: None to good.

9.3. Compromise Protection

Compromise protection is a collective term for a number of security services. These services, described below, are all concerned with the secrecy, or non-disclosure of information transfer between peer entities through the computer communications network. Physical security, such as protected wireways, can also provide transmission security. The network manager or sponsor must decide on the balance between physical, administrative, and technical security. This document only addresses technical security.

9.3.1. Data Confidentiality

+ Functionality

Data confidentiality service protects data against unauthorized disclosure. Data confidentiality is mainly compromised by passive wiretapping attacks. Passive attacks consist of observation of information passing on a link. Release of message content to unauthorized users is the fundamental compromise.

Prevention of release of message contents can be accomplished by applying an encryption mechanism. (See also the Encryption Mechanism section.) The granularity of key distribution is a trade-off between convenience and protection. Fine granularity would employ a unique key for each sensitivity level for each session; coarse granularity would employ the same key for all sessions during a time period.

The network must provide protection of data from unauthorized disclosure. Confidentiality can have the following features:

1. Confidentiality of all user-data on a specific protocol layer connection. Note: depending on use and layer, it may not be appropriate to protect all data, e.g., expedited data or data in a connection request.
2. Confidentiality of all user-data in a single connectionless datagram
3. Confidentiality of selected fields within the user-data of an PDU

Basis for Rating: Presence or absence of each feature.

Evaluation Range: None or present.

+ Strength of Mechanism

Physical protection and encryption are the fundamental techniques for protecting data from compromise. Through their use, release of message content and traffic analysis can be prevented.

Basis for Rating: The evaluation of data confidentiality mechanisms is outside the scope of this document. The cognizant authorities will evaluate the mechanisms relative to a specific environment according to their own rules and procedures.

Evaluation Range: Sensitivity level of data approved to protect.

+ Assurance

Basis of rating: Assurance is concerned with guaranteeing or providing confidence that features to address Data Confidentiality threats have been implemented correctly and that the control objectives of each feature have been actually and accurately achieved. Blacker is an example of such an application of a TCB for high assurance of data confidentiality.

Many of the assurance approaches for data confidentiality are common to other security services. See the General Assurance section for further information.

Evaluation Range: None to good.

9.3.2. Traffic Flow Confidentiality

+ Functionality

Traffic flow confidentiality service protects data against unauthorized disclosure. Traffic analysis is a compromise in which analysis of message length, frequency, and protocol components (such as addresses) results in information disclosure through inference.

Traffic flow confidentiality is concerned with masking the frequency, length, and origin-destination patterns of communications between protocol entities. Encryption can effectively and efficiently restrict disclosure above the transport layer; that is, it can conceal the process and application but not the host computer node.

The OSI Addendum- notes: "Traffic padding mechanisms can be used to provide various levels of protection against traffic analysis. This mechanism can be effective only if the traffic padding is protected by a confidentiality

- ISO 7498/Part 2 - Security Architecture, ISO / TC

97 / SC 21 / N1528 / WG 1 Ad hoc group on Security, Project 97.21.18, September 1986.

service."

Basis for Rating: Presence or absence.

Evaluation Range: None or present.

+ Strength of Mechanism

Physical protection, encryption, and traffic padding are the fundamental countermeasures for traffic analysis.

Basis for Rating: The evaluation of traffic confidentiality mechanisms are outside the scope of this document. The cognizant authorities will evaluate the mechanisms relative to a specific environment according to their own rules and procedures.

Evaluation Range: Sensitivity level of data approved to protect.

+ Assurance

Basis for rating: Assurance is concerned with guaranteeing or providing confidence that features to address Traffic Confidentiality threats have been implemented correctly and that the control objectives of each feature have been actually and accurately achieved.

Many of the assurance approaches for traffic confidentiality are common to other security services. See the General Assurance section for further information.

Evaluation Range: None to good.

9.3.3. Selective Routing

+ Functionality

"Routing control is the application of rules during the process of routing so as to choose or avoid specific networks, links or relays-.... Routes can be chosen either dynamically or by prearrangement so as to use only physically secure sub-networks, relays, or links. End-systems may, on detection of persistent manipulation attacks, wish to instruct the network service provider to establish a connection via a different route. Data carrying labels may be forbidden by the security policy to pass through certain sub-networks, relays or links. Also the initiator of a connection (or the sender of a connectionless data unit) may specify routing caveats requesting that specific sub-networks, links or relays be avoided."

- ISO 7498/Part 2 - Security Architecture, ISO / TC

97 / SC 21 / N1528 / WG 1 Ad hoc group on Security, Project 97.21.18, September 1986.

For example, there are national laws and network administration policies governing individual privacy rights, encryption, and trans-border data flow. A user in a end system may wish to specify countries through which certain information should not flow.

Basis for Rating: Presence or absence.

Evaluation Range: None or present.

+ Strength of Mechanism

Basis for Rating: The factors discussed under Supportive Primitives (Section 7) apply.

Evaluation Range: None to good.

+ Assurance

Basis for Rating: The General Assurance Approaches apply.

Evaluation Range: None to good.

Table II-1. Part II Assurance Rating Relationship to Part I Evaluation

(note: Table not included)

Table II-2. Evaluation Structure for the Network Security Services

(note: Table not included)

Appendix A

Evaluation of Network Components

A.1. Purpose

Part I of this Trusted Network Interpretation (TNI) provides interpretations of the Trusted Computer Security Evaluation Criteria (TCSEC) appropriate for evaluating a network of computer and communication devices as a single

system with a single Trusted Computing Base (TCB), called the Network Trusted Computing Base (NTCB), which is physically and logically partitioned among the components of the network. These interpretations stem from the recognition that networks form an important and recognizable subclass of ADP systems with distinctive technical characteristics that allow tailored interpretations of the TCSEC to be formulated for them.

An extension of this view of networks can be taken: that a trusted network represents a composition of trusted components. This view is sound, consistent with the first, and useful. The approach to evaluation of a network suggested by this view is to partition the system into components, rate each component to determine its security-relevant characteristics, and then evaluate the composition of the components to arrive at an overall rating class for the network. This approach aids in the assigning of an overall network evaluation class in two ways: 1) it allows for the evaluation of components which in and of themselves do not support all the policies required by the TCSEC (which will then contribute to the overall evaluation of any network which uses the evaluated component), and 2) it allows for the reuse of the evaluated component in different networks without the need for a re-evaluation of the component.

This approach to evaluation does not negate or override any of the interpretations presented in Part I of this document, which describe the global characteristics of a trusted network. In order to present a unified and self-consistent exposition within Part I of the document, a deliberate choice was made to express the basic network interpretations in terms of the view that networks are instances of ADP systems to which the TCSEC are applied on a system-wide basis. This choice allows Part I to follow the TCSEC closely because the basic structural model underlying the TCSEC, that of a system with a single Trusted Computer Base (TCB), has not been altered.

This appendix provides guidance for the evaluation of the individual components of a trusted network. The component evaluation guidelines constitute a refinement and application of the total network interpretations expressed within Part I and Part II of this document, and are intended to support the eventual evaluation of a network or network subsystem product to attain an overall network class using the Part I interpretations. Note that Part II applies to components without further interpretation. No implication is intended in this appendix that all networks must be composed from evaluated components: it is conceivable that a complete network could be evaluated as a whole using the system interpretations presented in Part I. In many practical cases, however, the techniques presented here for considering first the individual components, and then their

composition into an evaluable whole, constitutes a viable and attractive means for actually conducting the evaluation of the system under Part I interpretations.

Three major issues must be confronted by the architect or evaluator of a trusted system when the partitioned viewpoint is applied:

1. How is the network to be partitioned in such a way that evaluation of individual components will support eventual evaluation of the entire network?
2. What evaluation criteria should be applied to each component when rating that component?
3. How can the composition of rated components be evaluated?

The first of these issues is addressed in the separate Appendix B, Rationale Behind NTCB Partitions. The remaining two issues are addressed in this Appendix: the first, in section A.1.1 and section A.3, and the last in section A.2.

Section A.1.1 presents a taxonomy (classification scheme) for processing components based upon subordinate policy elements to be enforced, as well as the rating structure for individual components.

Section A.2 presents techniques and guidelines for the composition of rated processing components to achieve particular system ratings for the assembled network. This guidance is based on characterizing each component according to the policy elements supported where these are organized into the four broad policy areas of Mandatory Access Controls, Discretionary Access Controls, Identification and Authentication, and Audit support.

Section A.3 presents specific evaluation guidance in terms of the network interpretations articulated in Part I of this document, to allow individual processing components to be rated preparatory to their utilization in a trusted network. The sections are organized according to component type, as defined in section A.1.1. For each component type, the applicable interpretations, from Part I, are provided, organized according to rating class.

A.1.1. Component Taxonomy and Rating Structure

The primary difference between a processing component, regarded as part of a larger network system, and regarded as a stand-alone ADP system is that as a stand-alone system all of the TCSEC requirements for a particular class must be met: for policy requirements (i.e., what features the sys-

tem must support) the intent of the TCSEC is to enforce a collection of features which are felt to be operationally complete and consistent for a total system. In the context of a larger system, however, it may well be (and usually is) the case that the set of policy-related features to be supported by the component need not be the complete set required for a stand-alone system: features not supplied by one component for the system are supplied by another.

In rating a product for potential use as a network component, we would like, in theory, to be able to characterize its security properties exactly: in practice, we shall be content to identify the component as being of a particular type (which identifies the general policy elements the component supports) and of a particular evaluation class (which identifies the assurance levels provided for each supported feature), and the target architecture. The description of the target architecture shall include a description of the services that must be provided by other devices.

In order to limit the number of component types we break the ``maximal'' set of policy-related features, defined by the TCSEC for A1 systems, into four relatively independent categories which can be characterized as supporting Mandatory Access Controls (MAC), Discretionary Access Controls (DAC), Audit, and Identification and Authentication. (In various tables and text in the remainder of this appendix, these categories will be given the one-letter designations M, D, A, and I, respectively).

A given component can be intended (by the component sponsor) or required (by the network sponsor) to provide any combination of M, D, A or I functionality. Logically, then, there are sixteen different component types which can be rated using the guidelines of section A.3, corresponding to the sixteen possible combinations of M, D, A, and I theoretically possible. Of these combinations one (no M, no D, no A, no I) typifies a component intended (or required) to enforce no security policy whatsoever, and therefore has no TCSEC requirements to meet and need not be evaluated. However, it is still possible to utilize such components as part of a secure network system, if the architecture of the system induces a nil subordinate policy upon the component. The remaining component types are denoted M, D, A, I, MD, MA, MI, DA, DI, IA, MDA, MDI, MIA, IAD, and MIAD with the obvious meanings (for example, an "MIA component" supports aspects of Mandatory, Audit, and Authentication and ID policies, with the exact features provided being specified in section A.3 depending upon both evaluation class and type).

Table A1. Component Type Maximum and Minimum Class

COMPONENT	TYPE	MIN CLASS	MAX CLASS
M		B1	A1

D	C1	C2+
I	C1	C2
A	C2	C2+
DI	C1	C2+
DA	C2	C2+
IA	C2	C2+
IAD	C2	C2+
MD	B1	A1
MA	B1	A1
MI	B1	A1
MDA	B1	A1
MDI	B1	A1
MIA	B1	A1
MIAD	B1	A1

In addition to a type based upon the policy elements supported, an evaluated processing component is assigned a single evaluation class. In order to achieve a particular class, the component must meet all of the guidelines for that rating level for the applicable component type provided in section A.3. In general, these guidelines are straightforward interpretations of the TCSEC for the subset of policy features to be provided. Each component type has a maximum and minimum class listed in Table A1 below. To achieve a particular class, a component must meet appropriate requirements for policy, assurance, accountability, and documentation.

The maximum class for each component type is derived from the TCSEC, and is that evaluation class which imposes the highest requirement relevant to the component type. Similarly, the lowest class available for each component type is the TCSEC evaluation class which first imposes a requirement relevant to that component type.

Exceptions to this general approach have been made for the requirements for DAC and Audit support at the B3 level as the additional support for these policy categories at these levels (namely, the provision of ACL's for DAC and for real-time alarms for Audit) are not at the high level of assurance provided for the B3 MAC support. It is considered more appropriate to use the notation of C2+ for component types including D or A, but not M which meet the functional requirements of the B3 system ratings for D or A.

Components including support for I may be required to provide identification and authentication support for DAC (at possibly relatively low levels of assurance) or both DAC and MAC (at relatively high levels of assurance). Therefore, rating levels ranging from C1 to A1 for type I components have been provided. The ratings above B2 reflect the need for added assurance for the label integrity for the

MAC label information, rather than any additional requirements for features.

Components including support for I are required to provide Identification and Authentication which supports the DAC Policy. The TCSEC Identification-Authentication requirements for establishing a user clearance are reflected in M Components, since this requirement is in essence establishing a security label for a user.

Components of multiple types have been given minimum and maximum levels based upon meaningful combinations of the included types.

It might be noted in passing that a C1 stand-alone system has exactly the same certification requirements as a C1 DI component, a C2 system as a C2 IAD component, and B1-A1 systems as B1-A1 MIAD components.

A.2. Composition Rules

A.2.1. Purpose

In order for a (sub)system composed of components to be assigned a rating, the components that make up the network must be interconnected in such a way that the connections do not violate the assumptions made at the time the components were individually evaluated. This section presents the rules for the composition of evaluated components to form an evaluatable (sub)system and the method for assigning a rating to a (sub)system conforming to the rules.

This section does not consider the relative risk of utilizing the evaluated (sub)system to separate data at various levels of sensitivity: that is the role of the environmental security requirements, such as those of Computer Security Requirements, Guidance for applying the Department of Defense Trusted Computer System Evaluation Criteria in Specific Environments, CSC-STD-003-85. This section presents a technical basis for assigning a rating to a (sub)system which is composed of more than one component. The rating assigned indicates a minimum level of security as provided by the rated (sub)system as a whole.

Components must provide interfaces to support the other policies as required.

The composition rules are divided up according to the

15 possible combinations of the four policies supported by evaluated components (i.e., Mandatory Access Control, Discretionary Access Control, Audit, and Identification-Authentication).

A.2.2. Discretionary Access Control (D-Only) Composition

Rules

The rules presented below are based on the concept of properly composing a new component from previously evaluated components. Specifically, the rules presented in this section deal with the composition of a component with respect to the DAC Policy of the Network. It is expected that the composition of a D-Component will require significant engineering and system architectural consideration.

When a D-Component is evaluated, it will be evaluated against some stated Network DAC Policy and a stated target Network Security Architecture. Included in the component definition will be a statement of supported protocol for passing Identifiers which will be used as the basis for making DAC decisions. The stated protocol will be evaluated to assure that it is sufficient to support the target Network DAC Policy (e.g., if the Network DAC Policy is that access be designated down to the granularity of a single user, then an Identification Protocol which maps all users into a single Network ID would not suffice).

The type of Components discussed below, D-Components, are components that have received a rating relative to DAC (e.g., C1-C2+ D-Only Component, C1-C2+ DI Component, B1-A1 MD Component, etc.). The rules of this section are concerned only with the composition of these components with respect to the DAC Policy.

A.2.2.1. Composition of Two D-Components

Whenever two D-Components are directly connected the Identification Passing Protocol used to pass identifiers from one component to the other (for the purposes of making DAC decisions) must be the same in both components. It must be the case that the Identification Passing Protocol provided by the composed component must support the Identification Passing Protocol of the target Network Architecture. In addition, the composed DAC Policy (defined by the combination of the DAC Policy provided by one component over the named objects under its control and by the DAC Policy provided by the other component over the named objects under its control) must be shown to be able to support the target Network DAC Policy.

A.2.2.2. Discretionary Access Control Policy Composition

Rating

Given that a component is composed as described above, the evaluation class assigned to the composed component, with relation to DAC, will be the rating of the lowest class assigned to any D-Component within the composed component.

A.2.3. Identification-Authentication (I-Only) Composition

Rules

The rules presented below are based on the concept of properly composing a component from evaluated components. Specifically, the rules presented in this section deal with the composition of a component with respect to the Identification and Authentication Policy of the Network. It is expected that the composition of an I-Component will require significant engineering and system architectural consideration.

When an I-Component is evaluated it will be evaluated against some stated Network Identification-Authentication Policy and a stated target Network architecture. Included in the component definition will be a statement of the supported protocol for communicating User Identification and Authentication Information, and the interfaces provided by the I-Component. The composition of two I-Components must maintain the protocol which supports the Identification-Authentication Policy of the Network. In addition the interfaces provided by the composed I-Component, which support the stated protocol, must be identified.

A.2.3.1. Identification-Authentication Composition Rating

Given that a component is composed as described above, the evaluation class assigned to the composed component, with relation to Identification-Authentication, will be the rating of the lowest class assigned to any I-Component within the composed component.

A.2.4. Audit (A-Only) Composition Rules

The rules presented below are based on the concept of properly composing a component from evaluated components. Specifically, the rules presented in this section deal with the composition of a component with respect to the Audit Policy of the Network. It is expected that the composition of an A-Component will require significant engineering and

system architectural consideration.

When an A-Component is evaluated it will be evaluated against some stated Network Audit Policy and a stated target Network architecture. Included in the component definition will be a statement of the supported protocol that the component uses for receiving Audit information. The composition of two A-Components must maintain the protocol which supports the Audit Policy of the Network.

A.2.4.1. Audit Composition Rating

Given that a component is composed as described above, the evaluation class assigned to the composed component, with relation to Audit, will be the rating of the lowest class assigned to any A-Component within the composed component.

A.2.5. Mandatory Access Control (M-Only) Composition Rules

The rules presented below are based on the concept of properly composing a component from two directly connected (at the physical layer) components. Specifically, the rules presented in this section deal with the composition of a component with respect to the MAC Policy of the Network.

The MAC Composition Rules provide a strong guarantee that if the network is composed of directly connected, evaluated components, and each connection meets the MAC Composition Rules, the Network MAC Policy will be supported. These rules permit the recursive definition of a component based on the MAC Policy.

The MAC Composition Rules are divided into two sections. The first section addresses the composition of a component from two directly connected components with multilevel devices at each end of the connection. The second section addresses the composition of a component from two directly connected components with single-level devices at each end of the connection.

The type of Components discussed below, M-Components, are components which have received a rating relative to the MAC Policy (e.g., B1-A1 M-Only Components, B1-A1 MD-Components, B1-A1 MI-Components, etc.).

A.2.5.1. Multilevel Devices

Whenever two M-Components are directly connected, via a communication channel, with a multilevel device at each end of the connection, the labeling protocol (as required by the

Exportation to Multilevel Devices requirements, sections 3.1.1.3.2.1, 3.2.1.3.2.1, 3.3.1.3.2.1, and 4.1.1.3.2.1) must be the same at the network interface to both devices.

Whenever two Class B1 M-Component are directly connected, the range of sensitivity labels denoted by the maximum and minimum levels (System High and System Low) associated with each of the Class B1 M-Components must be the same. (This is because there are no explicit device labels for Class B1.)

Whenever a Class B1 M-Component is directly connected to a Class B2-A1 M-Component, the range of sensitivity labels denoted by the maximum and minimum levels (System High and System Low) associated with the Class B1 M-Component must be the same as the range of sensitivity labels denoted by the maximum and minimum levels associated with the multilevel device of the Class B2-A1 M-Component.

Whenever two Class B2-A1 M-Components are directly connected with a multilevel device at each end of the connection, the range of sensitivity labels denoted by the maximum and minimum levels associated with the each of the connected devices must be the same.

A.2.5.2. Single-Level Devices

Whenever two M-Components are directly connected with a single-level device at each end of the connection, the sensitivity level associated with the two devices must be the same.

Whenever two Non-M-Components are directly connected the maximum sensitivity level of data processed by the two Non-M-Components must be the same.

A.2.5.3. Mandatory Access Control Policy Composition Rating

Given that a component is composed as described in sections 2.5.1 and 2.5.2 above, the evaluation class assigned to the composed component, with regard to MAC, will be the rating of the lowest class assigned to any M-Component within the composed component.

A.2.6. DI-Component (D-Only and I-Only) Composition Rules

The rules presented below are based on the concept of properly composing a component from evaluated components. Specifically, the rules presented in this section deal with the composition of a component with respect to the DAC Policy and the Identification-Authentication Policy of the Net-

work. It is expected that the composition of a DI-Component will require significant engineering and system architectural consideration.

Whenever an I-Component and a D-Component are composed to form a DI-Component the DI-Component must preserve the Network DAC Policy of the D-Component. This implies that, depending on the DAC Policy, a protocol for receiving DAC information and returning data might be required for each DAC interface. This protocol must be able to support the Network DAC policy. (Note that if the Network DAC policy is defined such that access decisions are based on the user being a "member of the network group", i.e., is a legitimate user of another component, then the DAC interface may not require any identifiers to be passed to the DI-Component.)

In addition, for class C2 and above, the composed DI-Component must preserve the Audit Interface(s) used for exporting audit information from the D-Component and the I-Component. This implies that the DI-Component must provide a means for exporting audit information generated by actions taken within each of its parts.

The DI-Component may provide Identification-Authentication support services to other components. In this case the Identification Interface of the DI-Component must be defined and a protocol established for this interface which is able to support the Network I/A Policy. In this case the DI-Component may be further composed with other D-Only Components to form new DI-Components, using the rules defined above.

However, it is not necessary that the DI-Component provide Identification-Authentication services to other components. In this case the DI-Component may only be composed with other components (i.e., DI-Components, MIAD-Components, MI-Components, etc.) which are also self sufficient with respect to Identification-Authentication services.

If the composed DI-Component supports directly connected users then the DI-Component must, minimally, meet all the requirements for a Class C1 Network System.

A.2.6.1. DI-Component Composition Rating

Given that a component is composed as described above, and that the I-Component has an evaluation class of C1, the evaluation class assigned to the composed DI-Component, will be C1.

Given that a component is composed as described above, and that the I-Component has an evaluation class of C2, the evaluation class assigned to the composed DI-Component, will

be equal to the evaluation class of the D-Component.

A.2.7. DA (D-Only and A-Only) Composition Rules

The rules presented below are based on the concept of properly composing a component from evaluated components. Specifically, the rules presented in this section deal with the composition of a component with respect to the DAC Policy and the Audit Policy of the Network. It is expected that the composition of a DA-Component will require significant engineering and system architectural consideration.

Whenever an A-Component and a D-Component are composed to form a DA-Component, the DA-Component must preserve the Network DAC Policy of the D-Component. This implies that, depending on the DAC Policy, a protocol for receiving DAC information and returning data, might be required for each DAC interface. This protocol must be able to support the Network DAC policy. (Note that if the Network DAC policy is defined such that access decisions are based on the user being a "member of the network group", i.e., is a legitimate user of another component, then the DAC interface may not require any identifiers to be passed to the DI-Component.)

The DA-Component may provide Audit support services to other components. In this case the Audit Interface of the DA-Component must be defined and a protocol established for this interface, which is able to support the Network Audit Policy. In this case the DA-Component may be further composed with other D-Only Components to form new DA-Components, using the rules defined above.

However, it is not necessary that the DA-Component provide Audit services to other components. In this case the DA-Component may only be composed with other components (i.e., DA-Components, MIAD-Components, MA-Components, etc.) that are also self sufficient with respect to Audit services.

A.2.7.1. DA-Component Composition Rating

Given that a component is composed as described above, and that the D-Component has an evaluation class of at least C2, the evaluation class assigned to the composed DA-Component, will be the rating of the lowest class assigned to either of the two components which make up the composed component.

A.2.8. IA (I-Only and A-Only) Composition Rules

The rules presented below are based on the concept of

properly composing a component from evaluated components. Specifically, the rules presented in this section deal with the composition of a component with respect to the Identification-Authentication Policy and the Audit Policy of the Network. It is expected that the composition of a IA-Component will require significant engineering and system architectural consideration.

Whenever an IA-Component is composed of an I-Component connected to an A-Component, the IA-Component must preserve both the Network Audit Interface and Protocol of the A-Component and the Network Identification-Authentication Interface and Protocol of the I-Component. This implies that the composed IA-Component must provide an Audit Interface as well as a Identification-Authentication Interface. A protocol, for receiving Audit data, must be defined for each Audit Interface. This protocol must be able to support the Network Audit Policy. In addition, a protocol, for receiving Identification-Authentication data and returning authenticated user-ids, must be defined for each Identification Interface. This protocol must be able to support the Network I/A policy.

A.2.8.1. IA-Component Composition Rating

Given that a component is composed as described above, and that the I-Component has an evaluation class of at least C2, the evaluation class assigned to the composed IA-Component, will be the rating of the A-Component.

A.2.9. MD (M-Only and D-Only) Composition Rules

The rules presented below are based on the concept of properly composing a component from evaluated components. Specifically, the rules presented in this section deal with the composition of a component with respect to the MAC Policy and the DAC Policy of the Network. It is expected that the composition of an MD-Component will require significant engineering and system architectural consideration.

Whenever an MD-Component is composed from an M-Component directly connected to a D-Component, the composition rules, with respect to the MAC Policy, are that the M-Component must only connect to the D-Component via a single-level device, and the sensitivity level of the device must be the same as the maximum sensitivity level of data processed by the D-Component. Any network interfaces provided by the MD-Component via direct connections to the D-Component must be at the level of the D-Component.

The composition rules, with respect to the DAC Policy, are that any network interfaces provided by the MD-Component

(including those which only involve direct connections to the M-Component) must support the Identification Passing Protocol used by the D-Component. (Note that if the Network DAC policy is defined such that access decisions are based on the user being a ``member of the network group'', i.e., is a legitimate user of another component, then the DAC interface may not require any identifiers to be passed to the DI-Component.)

In addition, the composed MD-Component must ensure that any external requests for access to data under the control of the composed component are subject to both the MAC and DAC Policies of the original M and D Components.

A.2.9.1. MD-Component Composition Rating

Given that a component is composed as described above, and that the D-Component has an evaluation class of C2, the evaluation class assigned to the composed MD-Component, will be either B1 (if the evaluation class of the M-Component is B1) or B2 (if the evaluation class of the M-Component is greater than B1).

Given that a component is composed as described above, and that the D-Component has an evaluation class of C2+, the evaluation class assigned to the composed MD-Component, will be equal to the evaluation class of the M-Component.

A.2.10. MI (M-Only and I-Only) Composition Rules

The rules presented below are based on the concept of properly composing a component from evaluated components. Specifically, the rules presented in this section deal with the composition of a component with respect to the MAC Policy and the Identification-Authentication Policy of the Network. It is expected that the composition of an MI-Component will require significant engineering and system architectural consideration.

Whenever an MI-Component is composed from an M-Component directly connected to an I-Component, the composition rules, with respect to the MAC Policy, are that the M-Component must only connect to the I-Component via a single-level device, and the sensitivity level of the device must be the same as the maximum sensitivity level of data processed by the I-Component. Any network interfaces provided by the MI-Component via direct connections to the I-Component must be at the level of the I-Component.

In addition, the composed MI-Component must preserve the Audit Interface(s) used for exporting audit information from the M-Component and the I-Component. This implies that

the MI-Component must provide a means for exporting audit information generated by actions taken within each of its parts.

The MI-Component may provide Identification-Authentication support services to other components. In this case the Identification Interface of the MI-Component must be defined and a protocol established for this interface, which is able to support the Network I/A Policy. In this case the MI-Component may be further composed with other M-Only Components to form new MI-Components, using the rules defined above.

However, it is not necessary that the MI-Component provide Identification-Authentication services to other components. In this case the MI-Component may only be composed with other components (i.e., MI-Components, MIAD-Components, DI-Components, etc.) that are also self sufficient with respect to Identification-Authentication services.

The composed MI-Component must assure that MAC Policy and the Identification-Authentication Policy of the Network is supported on any direct User connections to the MI-Component. This implies that if the M-Component supports direct User connections, the M-Component must support a protocol on these connections such that Identification-Authentication information may be exchanged (with the I-Component) which will fully support the IA Policy of the Network.

A.2.10.1. MI-Component Composition Rating

Given that a component is composed as described above, and that the I-Component has an evaluation class of C2, the evaluation class assigned to the composed MI-Component will be equal to the evaluation class of the M-Component.

A.2.11. MA (M-Only and A-Only) Composition Rules

The rules presented below are based on the concept of properly composing a component from evaluated components. Specifically, the rules presented in this section deal with the composition of a component with respect to the MAC Policy and the Audit Policy of the Network. It is expected that the composition of an MA-Component will require significant engineering and system architectural consideration.

Whenever an MA-Component is composed from an M-Component directly connected to an A-Component, the composition rules, with respect to the MAC Policy, are that the M-Component must only connect to the A-Component via a single-level device and the sensitivity level of the device

must be the same as the maximum sensitivity level of data processed by the A-Component (probably Network High). Any network interfaces provided by the MA-Component via direct connections to the A-Component must be at the level of the A-Component.

The MA-Component may provide Audit support services to other components. In this case the Audit Interface of the MA-Component must be defined and a protocol established for this interface which is able to support the Network Audit Policy. In this case the MA-Component may be further composed with other M-Only Components to form new MA-Components, using the rules defined above.

However, it is not necessary that the MA-Component provide Audit services to other components. In this case the MA-Component may only be composed with other components (i.e., MA-Components, MIAD-Components, DA-Components, etc.) which are also self sufficient with respect to Audit services.

A.2.11.1. MA-Component Composition Rating

Given that a component is composed as described above, and that the A-Component has an evaluation class of C2, the evaluation class assigned to the composed MA-Component, will be either B1 (if the evaluation class of the M-Component is B1) or B2 (if the evaluation class of the M-Component is greater than B1).

Given that a component is composed as described above, and that the A-Component has an evaluation class of C2+, the evaluation class assigned to the composed MA-Component will be equal to the evaluation class of the M-Component.

A.2.12. IAD Composition Rules

The rules presented below are based on the concept of properly composing a component from evaluated components. Specifically, the rules presented in this section deal with the composition of a component with respect to the DAC Policy, the Identification-Authentication Policy, and the Audit Policy of the Network. It is expected that the composition of a IAD-Component will require significant engineering and system architectural consideration.

Whenever a IAD-Component is composed from directly connected components, the IAD-Component must conform to the composition rules for a DI-Component, a DA-Component, and an IA-Component. If the IAD-Component supports directly connected users then the IAD-Component must, minimally, meet all the requirements for a Class C2 Network System.

A.2.12.1. IAD-Component Composition Rating

Given that a component is composed as described above, and that the I-Component and D-Component each have an evaluation class of C2, the evaluation class assigned to the composed IAD-Component will be C2.

Given that a component is composed as described above, and that the I-Component has an evaluation class of C2 and the D-Component has an evaluation class of C2+, the evaluation class assigned to the composed IAD-Component will be the evaluation class of the A-Component.

A.2.13. MDA Composition Rules

The rules presented below are based on the concept of properly composing a component from evaluated components. Specifically, the rules presented in this section deal with the composition of a component with respect to the MAC Policy, the DAC Policy, and the Audit Policy of the Network. It is expected that the composition of a MDA-Component will require significant engineering and system architectural consideration.

Whenever a MDA-Component is composed from directly connected components, the MDA-Component must conform to the composition rules for an MD-Component, an MA-Component, and a DA-Component.

A.2.13.1. MDA-Component Composition Rating

Given that a component is composed as described above, and that the A-Component has an evaluation class of C2, and the D-Component has an evaluation class of C2 or higher, the evaluation class assigned to the composed MDA-Component will be either B1 (if the evaluation class of the M-Component is B1) or B2 (if the evaluation class of the M-Component is greater than B1).

Given that a component is composed as described above, and that the D-Component and A-Component each have an evaluation class of C2+, the evaluation class assigned to the composed MDA-Component will be equal to the evaluation class of the M-Component.

A.2.14. MDI Composition Rules

The rules presented below are based on the concept of properly composing a component from evaluated components.

Specifically, the rules presented in this section deal with the composition of a component with respect to the MAC Policy, the DAC Policy, and the Identification-Authentication Policy of the Network. It is expected that the composition of a MDI-Component will require significant engineering and system architectural consideration.

Whenever a MDI-Component is composed from directly connected components, the MDI-Component must conform to the composition rules for an MD-Component, an MI-Component, and a DI-Component.

A.2.14.1. MDI-Component Composition Rating

Given that a component is composed as described above, and that the I-Component and the D-Component each have an evaluation class of C2, the evaluation class assigned to the composed MDA-Component will be either B1 (if the evaluation class of the M-Component is B1) or B2 (if the evaluation class of the M-Component is greater than B1).

Given that a component is composed as described above, and that the I-Component has an evaluation class of C2, and the D-Component has an evaluation class of C2+, the evaluation class assigned to the composed MDI-Component will be equal to the evaluation class of the M-Component.

A.2.15. MIA Composition Rules

The rules presented below are based on the concept of properly composing a component from evaluated components. Specifically, the rules presented in this section deal with the composition of a component with respect to the MAC Policy, the Identification-Authentication Policy, and the Audit Policy of the Network. It is expected that the composition of a MIA-Component will require significant engineering and system architectural consideration.

Whenever a MIA-Component is composed from directly connected components, the MIA-Component must conform to the composition rules for an MI-Component, an MA-Component, and a IA-Component.

A.2.15.1. MIA-Component Composition Rating

Given that a component is composed as described above, and that the I-Component and the A-Component each have an evaluation class of C2, the evaluation class assigned to the composed MIA-Component, will be either B1 (if the evaluation class of the M-Component is B1) or B2 (if the evaluation class of the M-Component is greater than B1).

Given that a component is composed as described above, and that the I-Component has an evaluation class of C2 and the A-Component has an evaluation class of C2+, the evaluation class assigned to the composed MIA-Component, will be equal to the evaluation class of the M-Component.

A.2.16. MIAD Composition Rules

The rules presented below are based on the concept of properly composing a component from evaluated components. Specifically, the rules presented in this section deal with the composition of a component with respect to the MAC Policy, the DAC Policy, the Identification-Authentication Policy, and the Audit Policy of the Network. It is expected that the composition of a MIA-Component will require significant engineering and system architectural consideration.

Whenever an MIAD-Component is composed from directly connected components, the MIAD-Component must conform to the composition rules for an MIA-Component, an MDA-Component, an MDI-Component, and a IAD-Component. If the MIAD-Component supports directly connected users then the MIAD-Component must, minimally, meet all the requirements for a Class B1 Network System.

A.2.16.1. MIAD-Component Composition Rating

Given that a component is composed as described above, and that the I-Component and the D-Component each have an evaluation class of C2, the evaluation class assigned to the composed MIAD-Component will be either B1 (if the evaluation class of the M-Component is B1) or B2 (if the evaluation class of the M-Component is greater than B1).

Given that a component is composed as described above, and that the I-Component has an evaluation class of C2, and the D-Component and the A-Component each have an evaluation class of C2+, the evaluation class assigned to the composed MIAD-Component will be equal to the evaluation class of the M-Component.

A.3. Guidelines for Specific Component Evaluation

A.3.1. Mandatory Only Components (M-Components)

Mandatory Only Components are components that provide network support of the MAC Policy as specified in the Network Interpretation of the DoD Trusted Computer System Evaluation TCSEC. M-Components do not include the mechan-

isms necessary to completely support any of the 3 other network policies (i.e., DAC, Identification-Authentication, and Audit) as defined in the Interpretation.

M-Components belong to one of four classes B1, B2, B3, and A1 (as defined by the requirements below).

M-Components are rated according to the highest level for which all the requirements of a given class are met.

A.3.1.1. Overall Interpretation

In the requirements referenced, TCB will be understood to refer to the NTCB Partition of the M-Component. Also any reference to audit for an M-Component will be interpreted to mean "the M-Component shall produce audit data about any auditable actions performed by the M-Component". In addition the M-Component shall contain a mechanism for making the audit data available to an audit collection component.

A.3.1.2. Generally Interpreted Requirements

The requirements listed in the Table A2 apply directly to M-Components as interpreted in Part I of this interpretation.

Table A2. M-Component Requirements That Can Be Applied

Without Further Interpretation

(NOTE: Table not included)

A.3.1.3. Specifically Interpreted Requirements

The following requirements require additional interpretation as indicated. -

A.3.1.3.1. Subject Sensitivity Labels

+ Criteria

(Class B2 - Section 3.2.1.3.3; Class B3 - Section 3.3.1.3.3; Class A1 - Section 4.1.1.3.3)

+ Interpretation

An M-Component need not support direct terminal input in which case this requirement is not applicable. Any M-Component which does support direct terminal input must meet

- For brevity, the following TCSEC sections contain pointers to the sections of Part I of the Network Interpretation being interpreted, instead of the actual requirements.

the requirement as stated.

+ Rationale

The only way that a user can change the current level of the session is to be directly connected to a component that supports the MAC Policy. If the user is directly connected to a component that does not support the MAC Policy then the user will always operate at the level of the component to which he is directly attached. If the user is directly connected to a M-Component then this M-Component must meet the requirements as stated. M-Components which may be part of the network which do not directly communicate with users need not support this requirement since the requirement will be met by the M-Component with which the user is directly communicating.

A.3.1.3.2. Trusted Path

+ Criteria

(Class B2 - Section 3.2.2.1.1; Class B3 -
Section 3.3.2.1.1; Class A1 - Section 4.1.2.1.1)

+ Interpretation

An M-Component need not support direct user input (e.g., the M-Component may not be attached to any user I/O devices such as terminals) in which case this requirement is not applicable. Any M-Component which does support direct communication with users must meet the requirement as stated. In addition, an M-Component with directly connected users must provide mechanisms which establish the clearance of users and associate that clearance with the users current session.

+ Rationale

Trusted Path is necessary in order to assure that the user is communicating with the TCB and only the TCB when security relevant activities are taking place (e.g., authenticate user, set current session security level). However,

Trusted Path does not address communications within the TCB, only communications between the user and the TCB. If, therefore, an M-Component does not support any direct user communication then the M-Component need not contain mechanisms for assuring direct TCB to user communications.

In the case where an M-Component does support direct user communication the Clearance of the user must be established by the M-Component. There are three possible means of providing this support: a) all direct user connections are via single-level channels, where the maximum level of the channel equals the minimum level of the channel, and physical access to the channel implies clearance to the level of the channel; in this case there may exist no security relevant activities so that the applicable trusted path requirements may be met by reason of the device labels alone, b) some direct user connections are via single-level channels, where the maximum level of the channel does not equal the minimum level of the channel, and physical access to the channel implies clearance to the maximum level of the channel, c) some direct user connections are via single-level channels, where the maximum level of the channel does not equal the minimum level of the channel, and the M-Component contains some internal mechanism for mapping the user clearance to the range on the channel. The first two options map the user clearance to the activities of the user through external means. The third option requires some internal mechanism. Such a mechanism might be a user id/password/clearance database maintained by the M-Component. Another acceptable mechanism might be a protocol and interface definition within the M-Component for obtaining such information (via a multilevel channel - the channel is multilevel because it is passing labels, i.e., the user clearance) from some other M-Component.

A.3.1.3.3. System Architecture

+ Criteria

(Class B1 - Section 3.1.3.1.1; Class B2 - Section 3.2.3.1.1; Class B3 - Section 3.3.3.1.1; Class A1 - Section 4.1.3.1.1)

+ Interpretation

An M-Component must meet the requirement as stated. In this interpretation the words "The user interface to the TCB shall be completely defined..." shall be interpreted to mean the interface between the reference monitor of the M-Component and the subjects external to the reference monitor shall be completely defined.

+ Rationale

The M-Component may not have a direct user interface but is expected to support subjects which are not part of the TCB. It is important that the interface between the TCB and subjects external to the TCB be completely defined. (Note that in such a case the subjects are always internal to the component, viz., are "internal subjects").

A.3.1.3.4. Covert Channel Analysis

+ Criteria

(Class B2 - Section 3.2.3.1.3; Class B3 - Section 3.3.3.1.3; Class A1 - Section 4.1.3.1.3)

+ Interpretation

An M-Component must meet the requirement as stated. In addition, if the analysis indicates that channels exist that need to be audited (according to the Covert Channel Analysis Guideline), the M-Component shall contain a mechanism for making audit data (related to possible use of covert channels) available outside of the M-Component (e.g., by passing the data to an audit collection component).

+ Rationale

If an M-Component contains covert channels that need to be audited the M-Component must produce the audit data such that auditing can be performed. Since all covert channels in the network occur in an M-Component, the M-Component must be the source of the audit record which records the possible use of the covert channel.

A.3.1.3.5. Security Testing

+ Criteria

(Class B1 - Section 3.1.3.2.1; Class B2 - Section 3.2.3.2.1; Class B3 - Section 3.3.3.2.1; Class A1 - Section 4.1.3.2.1)

+ Interpretation

An M-Component must meet the requirement as stated except for the words "normally denied under the ... discretionary security policy," which are not applicable to an M-Component.

+ Rationale

An M-Component does not support a discretionary secu-

rity policy, and therefore testing for violations of such a policy is of no value.

A.3.1.3.6. Design Specification and Verification

+ Criteria

(Class B1 - Section 3.1.3.2.2; Class B2 - Section 3.2.3.2.2; Class B3 - Section 3.3.3.2.2; Class A1 - Section 4.1.3.2.2)

+ Interpretation

An M-Component must meet the requirement as stated.

Security Policy is interpreted to mean the MAC Policy supported by the component. Model is interpreted to be those portions of a reference monitor model that are relevant to the MAC Policy supported by the Component (e.g., the representation of the current access set and the sensitivity labels of subjects and objects, and the Simple Security and Confinement Properties of the Bell and LaPadula Model).

A.3.1.3.7. Trusted Facility Manual

+ Criteria

(Class B1 - Section 3.1.4.2; Class B2 - Section 3.2.4.2; Class B3 - Section 3.3.4.2; Class A1 - Section 4.1.4.2)

+ Interpretation

An M-Component must meet the requirement as stated except for the words ``The procedures for examining and maintaining the audit files as well as...'. These words are interpreted to mean "the mechanisms and protocols associated with exporting of audit data must be defined." Also, the words "...to include changing the security characteristics of a user", shall not be applicable to an M-Component.

+ Rationale

An M-Component does not maintain the audit files nor does it provide mechanisms for examining them. It must, however provide mechanisms for exporting the audit files and these mechanisms need to be defined in the Trusted Facility Manual. The M-Component also does not maintain user information.

A.3.1.4. Representative Application of M-Components

As an example of an M-Component, consider a MLS packet switch that provides MAC through a verified Security Kernel, as shown in Figure A1. This component supports 16 levels and 64 categories for non-discretionary access classes. The MLS packet switch is rated as an A1 M-Component against the requirements described above.

Such an A1 M-Component may, as an example, be used in a network as a Multilevel Packet Switch. The M-Component could be configured with several single-level channels and some number of multilevel channels. As part of the example, assume that the multilevel channels each have a maximum of Top Secret and a minimum of Secret. Also imagine that the single-level channels are either Top Secret or Secret. The Multilevel channels are directly connected to B2 hosts each with a system high of Top Secret and a system low of Secret. The single-level channels are directly connected to C2 hosts with some of them running at dedicated Secret and some running at dedicated Top Secret. One of the Dedicated Top Secret Hosts and one of the Dedicated Secret Hosts would be responsible for collecting the audit messages sent from the M-Component. In this fashion one could create a network which could permit the multilevel hosts to securely communicate with each other as well as with the single-level hosts. The separation necessary for such communications would be provided by the M-Component being used as a Multilevel Secure Packet Switch. It is noted that the composition rules of Section A.2 (A.2.5 in particular) result in an evaluation class of B2 for the overall NTCB.

A.3.2. Discretionary Only Components (D-Components)

Discretionary Only Components are components that provide network support of the DAC Policy as specified in the Network Interpretation of the DoD Trusted Computer System Evaluation TCSEC. D-Components do not include the mechanisms necessary to completely support any of the three other network policies (i.e., MAC, Identification-Authentication, and Audit) as defined in the Interpretation.

D-Components belong to one of three classes, C1, C2, and C2+ (as defined by the requirements below).

D-Components are rated according to the highest level for which all the requirements of a given class are met.

A.3.3. Overall Interpretation

In the requirements referenced, TCB will be understood to refer to the NTCB Partition of the D-Component. Also any reference to audit for an D-Component will be interpreted to mean "the D-Component shall produce audit data about any auditable actions performed by the D-Component." In addition the D-Component shall contain a mechanism for making the audit data available to an audit collection component.

A.3.3.1. Generally Interpreted Requirements

The requirements listed in Table A3 apply directly to D-Components as interpreted in Part I of this interpretation.

A.3.3.2. Specifically Interpreted Requirements

The following requirements require additional interpretation as indicated. -

- For brevity, the following TCSEC sections contain pointers to the sections of Part I of the Network Interpretation being interpreted, instead of the actual requirements.

A.3.3.2.1. Trusted Facility Manual

+ Criteria

(Class C1 - Section 2.1.4.2; Class C2 - Section 2.2.4.2; Class C2+ - Section 2.2.4.2)

+ Interpretation

A D-Component must meet the requirement as stated except for the words ``The procedures for examining and maintaining the audit files as well as...''. These words are interpreted to mean "the mechanisms and protocols associated with exporting of audit data must be defined."

+ Rationale

A D-Component does not maintain the audit files, nor does it provide mechanisms for examining them. It must,

however provide mechanisms for exporting audit data to an audit component, and these mechanisms need to be defined in the Trusted Facility Manual.

A.3.3.2.2. Design Documentation

+ Criteria

(Class C1 - Section 2.1.4.4; Class C2 - Section 2.2.4.4; Class C2+ - Section 2.2.4.4)

+ Interpretation

A D-Component must meet the requirement as stated. In addition the Design Documentation must include a description of the protocol used by the D-Component to communicate Subject permissions (i.e., user ids), where applicable, with other components. This protocol must be shown to be sufficient to support the DAC policy enforced by the D-Component.

+ Rationale

A D-Component does not maintain the user Identification-Authentication information. It may, however, use some form of authenticated user identification as a basis for making DAC decisions. Such information must be provided to the D-Component through the identification protocol. The protocol used by the D-Component may vary, but it must be shown to be adequate to support the DAC policy supported by the D-Component. As an example consider a simple DAC policy in which access is granted, or denied, on a per host basis. In this case the protocol used might be to statically assign a host-id to each port. All requests coming in from a given port would be associated with the access permissions allowed for that host. This protocol would not be adequate to support a DAC policy of access granted, or denied, on a per user basis.)

A.3.3.3. Representative Application of D-Components

As an example of a D-Component, consider a system that provides DAC through Access Control Lists on files, as shown in Figure A2. The system is rated as a C2+ D-Component against the requirements described above.

Figure A2. Representative Application of D-Components

B1: box wid 1.15i ht .95i "Class B3" "Host" B2: box wid 1.15i ht .96i "Class C2+" "Host" "(Network Audit" "Collec-

tion" "Facility)" at B1.e +(3i,0) B3: box wid 1.15i ht .96i "Class C2+" "D-Component" "(Single Level" "File Server" at B2.s -(1.75i,.30i) B4: box wid 1.85i ht .96i "Class C2" "Host" "(Network Identification &" "Authentication" "Facility)" at B1.s -(0,1.05i) B5: box wid 1.15i ht .96i "Class A1" "Host" at B2.s -(0,1.05i) B6: box invis "(S)" at B1.e +(.50i,.30i) B7: box invis "(S)" at B2.w -(0.50i, -.30i) B8: box invis "(S)" at B4.e +(.50i,.2) B9: box invis "(S)" at B5.w -(0.50i,.2) A1: arrow <-> from B4.e to (B3.s.x-(B3.s.x+.2,B5.w.y) to (B3.s.x+.2,B3.s.y) A3: arrow left 1i from B6.c +(.48,-.15i) A4: arrow right 1i from B7.c -(0.48i,.15i) A5: arrow down .39i from A4.w A6: arrow down

Such a C2+ D-Component may, as an example, be used in a network as a Single Level File Server. The D-Component could be configured with several communication channels (each of which would be connected to single-level devices with the same access class). As part of the example, consider all files on the system to be secret and all channels leaving the system to be connected to other single-level secret components or, in the case of multi-level components, to be connected to single-level secret devices. The documentation associated with the D-Component must specify the protocol used to pass user-ids and filenames. This protocol must be followed on each connection to the component. In addition the documentation must specify the protocol used to output audit information. The audit protocol must be exactly the same as the protocol of the audit node to which it is attached. It is noted that the composition rules of Section A.2 result in an evaluation class of B3 for the overall NTCB.

A.3.4. Identification-Authentication Only Components (I-

 Components)

Identification-Authentication Only Components are components that provide network support of the Identification-Authentication Policy as specified in the Network Interpretation of the DoD Trusted Computer System Evaluation TCSEC. I-Components do not include the mechanisms necessary to completely support any of the three other network policies (i.e., MAC, DAC, and Audit) as defined in the Interpretation.

I-Components belong to one of two classes, C1 and C2 (as defined by the requirements below).

I-Components are rated according to the highest level for which all the requirements of a given class are met.

A.3.4.1. Overall Interpretation

In the requirements referenced, TCB will be understood to refer to the NTCB Partition of the I-Component. Also any reference to audit for an I-Component will be interpreted to mean "the I-Component shall produce audit data about any auditable actions performed by the I-Component.'" In addition the I-Component shall contain a mechanism for making the audit data available to an audit collection component.

A.3.4.2. Generally Interpreted Requirements

The requirements listed in Table A4 apply directly to I-Components as interpreted in Part I of this interpretation.

Table A4. I-Component Requirements That Can Be Applied

Without Further Interpretation

(Note: Table not included)

A.3.4.3. Specifically Interpreted Requirements

The following requirements require additional interpretation as indicated. -

A.3.4.3.1. Trusted Facility Manual

- For brevity, the following TCSEC sections contain pointers to the sections of Part I of the Network Interpretation being interpreted, instead of the actual requirements.

+ Criteria

(Class C1 - Section 2.1.4.2; Class C2 - Section 2.2.4.2;
Class C2+ - Section 2.2.4.2)

+ Interpretation

An I-Component must meet the requirement as stated except for the words ``The procedures for examining and maintaining the audit files as well as...''. These words are

interpreted to mean "the mechanisms and protocols associated with exporting of audit data must be defined."

+ Rationale

An I-Component does not maintain the audit files, nor does it provide mechanisms for examining them. It must, however, provide mechanisms for exporting audit data to an audit component, and these mechanisms need to be defined in the Trusted Facility Manual.

A.3.4.3.2. Design Documentation

+ Criteria

(Class C1 - Section 2.1.4.4; Class C2 - Section 2.2.4.4; Class C2+ - Section 2.2.4.4)

+ Interpretation

An I-Component must meet the requirement as stated. In addition the Design Documentation must include a description of the protocol used by the I-Component to export Authenticated Subject Identifiers to other components.

+ Rationale

The Authenticated Identifiers provided by an I-Component will not be primarily used on the I-Component itself but instead will be used by other Components enforcing the network DAC policy. It is therefore necessary for the I-Component to define the protocol which it will use to pass authenticated user-ids to other components.

A.3.4.4. Representative Application of I-Components

As an example of an I-Component, consider a system which provides Identification and Authentication facilities, such as a TAC with a name server, as shown in Figure A3. The system is rated as a C2 I-Component against the requirements described above. The I-Component could be configured with several communication channels (each of which would be connected to single-level devices with the same access class). As part of the example, consider the TAC to be an unclassified TAC (i.e., accessible through the phone system without any encryption support) and all channels leaving the system to be connected to other single-level unclassified components or, in the case of multi-level components, to be connected to single-level unclassified devices. All authentication is done in the TAC, and Authenticated Ids are

passed to the other nodes of the network to be used as a basis for DAC decisions and audit entries. The documentation associated with the I-Component must specify the protocol used to pass user-ids to the attached components. This protocol must be supported on each connection to the component. In addition the documentation must specify the protocol used to output audit information. The audit protocol must be exactly the same as the protocol of the audit component to which it is attached. It is noted that the composition rules of Section 3 result in an evaluation class of B3 for the overall NTCB.

A.3.5. Audit Only Components (A-Components)

Audit Only Components are components which provide network support of the Audit Policy as specified in the Network Interpretation of the DoD Trusted Computer System Evaluation TCSEC. A-Components do not include the mechanisms necessary to completely support any of the three other network policies (i.e., MAC, DAC, and Identification-Authentication) as defined in the Interpretation.

A-Components belong to one of two classes C2 and C2+ (as defined by the requirements below). (The difference between a C2 A-Component and a C2+ A-Component is the support of real time alarms required by class B3 Audit.)

A-Components are rated according to the highest level for which all the requirements of a given class are met.

A.3.5.1. Overall Interpretation

In the requirements referenced, TCB will be understood to refer to the NTCB Partition of the A-Component.

A.3.5.2. Generally Interpreted Requirements

The requirements listed in Table A5 apply directly to A-Components as interpreted in Part I of this interpretation.

A.3.5.3. Specifically Interpreted Requirements

The following requirements require additional interpretation as indicated. -

- For brevity, the following TCSEC sections contain pointers to the sections of Part I of the TNI being interpreted, instead of the actual requirements.

A.3.5.3.1. Design Documentation

+ Criteria

(Class C2 - Section 2.2.4.4; Class C2+ - Section 2.2.4.4)

+ Interpretation

An A-Component must meet the requirement as stated. In addition the Design Documentation must include a description of the protocol used by the A-Component to import Audit Data from other nodes.

+ Rationale

The Audit component will potentially be used for collection of audit data generated on many different components. Each of these components must be able to transfer the information to the A-component in a form that will allow the A-Component to create an audit record. The mechanism for defining the acceptable form of information is the protocol used by the audit component.

A.3.5.4. Representative Application of A-Components

As an example of an A-Component, consider a system that provides Audit Collection and review facilities for a network environment, as illustrated in Figure A4. The system is rate C2+ against the requirements described above.

As part of the example, consider the A-Component to be operating at System High (Top Secret) collecting information from several components through single-level (Top Secret) channels. The A-Component provides auditing functions for the network as a whole. The A-Component defines an audit protocol which is used by each of the components to communicate information to the A-Component which results in the creation of audit records. Note that in this example the Auditor (i.e., the person responsible for reviewing audit files) is accessing the A-Component through an operators console attached to the A-Component. In a different scenario, it might be the case that the Auditor accesses the A-Component via another component, in which case the A-Component would be responsible for enforcing an access control policy that defined which users (i.e., the auditor) could view audit data. This would require the A-Component to establish a user-id passing protocol much like a D-Component. It is noted that the composition rules of Sec-

tion 3 result in an evaluation class of B3 for the overall NTCB.

Figure A1. Representative Application of M-Components

(Note: Figure not included)

Table A3. D-Component Requirements That Can Be Applied
Without Further Interpretation

(Note: Table not included)

Figure A3. Representative Application of I-Component

(Note: Figure not included)

Table A5. Audit Component Requirements That Can Be Applied
Without Further Interpretation

(Note: Table not included)

Appendix B

Rationale Behind NTCB Partitions

B.1. Purpose

Part I of this Trusted Network Interpretation (TNI) provides interpretations of the Trusted Computer Security Evaluation Criteria (TCSEC) appropriate for evaluating a network of computer and communication devices as a single system with a single Trusted Computing Base (TCB), called the Network Trusted Computing Base (NTCB), which is physically and logically partitioned among the components of the

network. Implicit to this approach is the view that the network to be evaluated (including the interconnected hosts) is analogous to a single stand-alone computer system, and can therefore be evaluated using the TCSEC under appropriate interpretation. It is the purpose of this appendix to provide the main technical rationale and illustrative examples supporting this view. This underlying rationale may also be of help to the sponsors and evaluators of networks and network components in understanding how a network can be cleanly partitioned into components in a way that will facilitate its eventual evaluation and certification. It is recognized that this appendix is in places quite theoretical and philosophical. Therefore, readers whose interest is primarily in applying the TNI without reviewing its derivation may choose not to study this appendix in detail.

The separate Appendix A, providing Interpretations for the Evaluation of Network Components, rests upon this view as well: the evaluation of particular network components is viewed as a useful preliminary step for the eventual evaluation of the network as a whole, which must proceed, however, in the context of an overall network architecture providing a clean decomposition of an overall network security policy into policies for the individual components. The overall architecture and design will, once individual component evaluations have been finished, support the final evaluation of the network as a sound composition of trusted elements, each enforcing its allocated policy, and together enforcing the policy defined for the entire network. Specific guidelines for actually partitioning the various network policy elements to components are presented under the relevant headings in the separate Appendix A for Evaluation of Network Components: the general rationale supporting the view that such a partitioning is possible is presented here.

It is emphasized that the view of what a network is (and how its NTCB may be partitioned into NTCB partitions completely contained in individual network components) described in this appendix is adopted with one goal in mind: the evaluation and assignment to the network of a single certification as meeting the TCSEC criteria for a given evaluation class. It is recognized that this goal may not be appropriate for every circumstance, or meet the needs of sponsors wishing to interconnect already existing systems. The risk assessment and accreditation of such systems is an important and interesting problem. It is not, however, the problem being addressed here, viz., the evaluation of an entire network which is to support a network security policy given a priori.

B.2. Background and Overview

B.2.1. Organization of this Appendix

The material within this appendix is organized as follows. Section B.3 discusses some considerations for properly formulating the policy to be enforced by the network NTCB, and its allocation to the various components of the network. Section B.4 presents an argument supporting the adequacy of the partitioned NTCB view and the conclusion that the reference monitor for an entire network may be implemented as a collection of locally autonomous reference monitors. Section B.5 discusses the idealization of intercomponent communications channels, assumed as an axiom in Section B.4, in the context of real communications channels, and provides insight into when the techniques of communications security, and when the techniques of trusted systems technology are applicable. Section B.6 provides additional rationale supporting the partitioned NTCB view.

B.3. Security Policy

The TCSEC Glossary defines ``Security Policy'' as ``the set of laws, rules, and practices that regulate how an organization manages, protects, and distributes sensitive information''. It should be noted that ``Security Policy'' is a distinct notion from that of ``Formal Security Policy Model'' and a ``Security Policy Model''. The ``Security Policy'' of an organization has the ultimate goal of controlling the access of people to information.

Because a Security Policy concerns, by definition, the access of people to sensitive information and includes both secrecy and integrity; it can, ideally, be stated in a manner that is free of computer, network, or communication jargon. Moreover, we would observe that the evaluation of a network ultimately is possible only if a single, uniform network security policy can be adopted by the organizations whose information is to be stored, transmitted, or processed by the network and its components. The existence of such a policy is a precondition of any attempt to evaluate a network or its components in the sense of this appendix. If a network is to be used to allow the sharing of information among many organizations, the definition of a mutually acceptable Security Policy applicable to that sharing must be an early goal during the design of the network if the successful certification of a network providing that capability is desired.

B.3.1. Mandatory Access Control Policies

One may observe that, for those access controls nor-

mally denoted as ``Mandatory Access Controls'', the definition of a mutually acceptable joint policy may be expected to be relatively straightforward, as such controls are based, by definition, upon the comparison of a label denoting the sensitivity of the information contained within an information repository with a user clearance denoting the formal authorization of a user to access that information. The definition of a jointly acceptable policy may involve the merging of several systems of classifications and clearances into a unified system; in practice, if the systems in use by the various organizations are not already identical, those responsible for the protection of information within each organization must determine which external user clearances will be honored as an adequate basis for providing access to which classes of information.

It may also be true that a particular organization may have no explicit policy describable as ``mandatory'' in the

sense defined by the TCSEC. (In particular, many commercial or private institutions may be so characterized)-. It is possible to formulate a trivial mandatory access control policy for such organizations, however, with a single access class and clearance level (i.e., every user belonging to the institution has clearance to access all information belonging to the institution, except as refined by less rigorous access controls). Thus, it could well be that an overall system of mandatory access controls, at the policy level, for an arbitrary collection of institutions wishing to share information using a network, can be resolved in a relatively straightforward way; at least in the sense that the policy issues and effects of particular decisions are easy to understand.

B.3.2. Discretionary Access Control Policies

Turning to those policies characterizable as involving Discretionary Access Controls, one finds substantially greater freedom in the sorts of policies an organization might adopt. The notion of ``Discretionary Access Controls'', as defined in the TCSEC Glossary, involves the restriction of access by users to information based upon the identity of the users or their membership in a particular group, as well as the ability of a user with authorized access to an object containing information to pass that authorization to other users or groups either directly, or

- See, for example, Steven B. Lipner, ``Non-Discretionary Controls for Commercial Applications'', IEEE Proceedings of the 1982 Symposium on Security and

Privacy, April 26-28, 1982, Oakland, CA.

indirectly (viz., by copying it and providing authorization to access the copy). Within these limits, there is an extremely broad range of permissible policies, differing in how users may be grouped, the sorts of named information repositories that may form the basis for access controls, the modes of access that may form the basis for controls, and the mechanisms that may be defined for users to limit or propagate permission to access information. One would expect, therefore, that when designing a network, the formulation of an overall Discretionary Policy by a group of organizations may require a period of intensive generalization of policy. Moreover, the overall policy resulting from this activity may be expected to depend, to a relatively large extent, upon the underlying capabilities and functionality ascribed to the network.

B.3.3. Supporting Policies

In addition to the basic access control policies (mandatory and discretionary) addressed by the TCSEC are additional capabilities relating to the accountability of individuals for their security-relevant actions. These capabilities are usually thought of as comprising ``supporting'' policy: they provide an environment that allows for the effective enforcement and monitoring of the basic access policies enforced.

Accountability requirements are comprised of two major policy subcategories: identification and authentication policy, and audit policy. The former supports both mandatory and discretionary access control policy by specifying the requirements for authenticating the identity and clearance of an individual prior to permitting access, is the basis for determining the clearance of an individual in the case of mandatory access policy, is the basis for determining the group membership of an individual in the case of discretionary access policy, and is the basis for recording the identity of the individual taking or causing an auditable action.

Audit policy proper provides for the recording of those security-relevant events that can be uniquely associated with an individual user, so that those responsible for the security of sensitive information may hold users accountable for the security-relevant actions they take.

The supporting policies adopted by different organizations may differ even more widely than discretionary access control policies. The task of formulating a mutually acceptable set of overall supporting policies may be expected to be even more challenging for the sponsors of a

network than for discretionary policy.

B.3.4. Formal Security Policy Model

As defined in the TCSEC, a Formal Security Policy Model has a mathematically precise statement of a Security Policy. Whereas the objective of stating Security Policy is to reflect the requirements imposed upon a system by external authority, the purpose of a Formal Security Policy Model is to serve as a precise starting point in the chain of arguments leading to the higher levels of assurance required for systems of the higher evaluation classes. Thus, the requirement to state a Formal Security Policy Model consistent with its axioms is first introduced at Evaluation Class B2; it is not introduced earlier because the chain of arguments needed for lower evaluation classes does not require mathematical precision at their onset. The point of this observation is that the definition of a Formal Security Model is not a gratuitous requirement, but serves the purpose of facilitating construction of the chain of arguments required for the higher evaluation Classes.

Current practice requires a formal security policy model only for the access control policies to be enforced. The model is a representation of the reference monitor for a specific class of systems. The choice of model representation is strongly influenced by the technical characteristics of the system to be built, as the feasibility and economy of constructing the chain of assurance arguments needed to support a class B2 or above evaluation is typically substantially increased by utilizing a model that has an intuitively attractive resemblance to the abstractions of subject, object, and access properties of the target system.

As previously described, the reference monitor for a partitioned NTCB is composed of a collection of security kernels for individual components. In order to obtain the required levels of assurance that each such security kernel works correctly, a Formal Security Policy Model must be formulated for each such component. We would argue, however, that it is too restrictive to require that the formal model for each security kernel be the same, or that an overall model be formulated for the network, provided that each model is shown by convincing arguments to correctly represent the overall Security Policy, as allocated to the component. As the only function of a formal model is to support the evaluation of a security kernel, the sponsors and designers of network components should be free to choose that model which will most efficiently serve this purpose, relative to the engineering characteristics of the component; subject, of course, to the requirement that the model be an accurate representation of the Security Policy to be enforced by the component.

B.3.5. Summary of Policy Considerations for a Network

In summary, a precondition for the evaluation of a networked system of computers is the formulation of overall mandatory (when applicable), discretionary, and supporting policies, mutually acceptable to the managements of the organizations involved, and stated in terms of people accessing information (i.e., free, to the extent feasible, of computer and network jargon). In the case of mandatory policy, we would expect the formulation of an overall policy to involve the relatively straightforward issues of how clearances in use by one organization are to relate to the information access classes in use by another organization: the formulation of appropriate discretionary and supporting policies may be expected to be more challenging and significantly influenced by the particular network architecture chosen.

B.4. Derivation of the Partitioned NTCB View

B.4.1. Introduction to the Partitioned NTCB Concept

Using the definitions provided above, the following conclusion may be stated: if it is supposed (1) that a subject is confined to a single component throughout its lifetime, (2) that it may directly access only objects within its component, (3) that every component contains a component reference monitor that mediates all accesses made locally (and enforce the same access control policy), and (4) that all communications channels linking components do not compromise the security of the information entrusted to them, one may conclude that the total collection of component reference monitors is a reference monitor for the network. The conclusion follows because (1) all network accesses are mediated (because there are no non-local accesses); and (2) the network reference monitor cannot be tampered with (because none of its component reference monitors can be tampered with), and it is simple enough to validate (its correct operation, under the suppositions given, is assured if the correct operation of each of its component reference monitors is individually assured - the stricture against access across components prevents the introduction of additional complexity).

It is useful, before expanding this basic argument within the context of non-idealized network systems, to examine briefly the individual preconditions (axioms) that must be met in order for the conclusion to be valid, and state concisely why there is reason to believe that each can be achieved within the current state of network technology.

Generally, the crucial step the sponsors and architects of a proposed network must perform (prior to its evaluation) is its partitioning into components and communication channels in such a way that all of the axioms can be easily validated by the evaluators.

The first axiom is that regarding confinement of subjects (to a single component). Adoption of the conventional notion of a subject as a <process, domain> pair is adequate to fulfill this axiom, provided it is recognized that limiting the access of subjects to objects within the same component ensures that no domain encompasses objects from more than one component. It follows that no subject may ``move'' from one component to another. Even if we permit (as is sometimes done) the notion of a ``remote process'', once execution begins in a remote component a new subject has been introduced (because there has been a change in protection domains).

The second axiom requires that a subject be able to directly access objects only within the component with which the subject is associated. The major theoretical issue to be confronted is to understand how information may be transmitted between components without the sharing of objects between them. This issue is explored in some depth in Section B.5. Logically, the connection of components by an ideal communication channel is viewed as involving the transfer of information from one device to another without the existence of an intermediate object. (i.e., information ``in motion'' is not regarded as an object - a view which seems valid provided that no subjects may access it until it ``comes to rest'' within the destination component - and is then within an object again). This view is consistent with the TCSEC Glossary definition of ``object'' which includes the sentence, ``Access to an object potentially implies access to the information it contains''. For a Security-Compliant communication channel (discussed in Section B.6.2), there are no subjects with potential access to the information being transmitted while it is in transit: it is

therefore unnecessary (and misleading) to treat such information as an object. (This argument is invalid for complex channels, which contain internal subjects, which is the reason that such channels must be further partitioned.)

The third axiom requires that every component contain a component reference monitor which enforces that part of the network access control policy relevant to subjects and objects within the component. In validating this axiom, it is important to understand that for certain components, a degenerate component reference monitor may suffice (e.g., for a dedicated component for which all subjects and objects have, by virtue of the system architecture, the same authorization and sensitivity, respectively, so that no local

access attempts need be denied on the basis of policy enforcement). It is logically equivalent, in such cases, to claim that there is a reference monitor (which never does anything) or that there is no reference monitor (because nothing ever needs to be done). It is also important to understand that each reference monitor need only to enforce that subset of access control policy relevant (in terms of the network system architecture) to the local accesses possible within the component.

The fourth axiom requires that communications channels between components not compromise the security of sensitive information entrusted to them. Establishing that this axiom is actually met is a complex problem with some issues dealt with during system evaluation, and others during accreditation for use. A detailed discussion of the issues involved is provided in Section B.6 of this Appendix. Until that discussion, the validity of the axiom is assumed as a boundary condition allowing the evaluation of individual trusted components of the network, and their composition into a complete system.

B.4.2. Overview of the Argument for a Partitioned NTCB

To present the concept of a partitioned NTCB, and show how the TCSEC Criteria may be applied to it, an application analogous to a network of ``loosely-coupled'' NTCB partitions is described as running upon a single, stand-alone computer system with a TCB assumed to be evaluatable in terms of the existing Criteria. A series of transformations is then performed upon the simulation, that convert it into the hypothesized network with a single, partitioned NTCB. This argument is meant to demonstrate that the notion of partitioning a single NTCB into a set of loosely-coupled NTCB partitions is conceptually sound and does not require a radical departure from current evaluation practice. In effect, the argument serves as a constructive proof (although informally stated) that a trusted network is simply an instance of a trusted computer system.

B.4.3. Characterization of the Target Monolithic System

Consider first a multiprocessor, multiprogrammed monolithic computer system, presumed to conform to the TCSEC Criteria at, for example, a Class B2 level or higher. It has a Formal Security Policy Model, e.g., the Bell and LaPadula model, and it has been shown that the system is a valid interpretation of that model. In the presumed system, suppose that the code and data of the TCB is shared among the concurrently executing processors, which are tightly-coupled on a single bus. (Worked examples of such systems targeted for Class B2 or higher exist). Since this is a monolithic

system with (it is presumed) an ``ordinary'' secure multiprogramming operating system, it can support a given process on any processor, which can (potentially) access any memory segments it may need to share with any other process on any other processor. Additionally, each process can use devices through I/O channels that are accessed by service calls to the TCB. In particular, assume that there are available multilevel I/O channels which can be controlled by multilevel trusted processes executing under the control of the TCB. Each multilevel channel conforms to the concept of a connected multilevel device as identified in the TCSEC Criteria.

B.4.4. Characterization of the Loosely-Coupled Trusted Network

Next, consider an arbitrary network architecture, consisting of various types of nodes (e.g., packet switches, network interface units, hosts, etc.) processing information at various levels, connected with communication channels, possibly multilevel. It is assumed that the network is secure, and meets the axioms described in section B.3.1, viz., (1) each subject is confined to a single component, (2) no subject may access an object within a different component, (3) each component possesses a locally autonomous reference monitor, (4) the communications channels are secure, in the sense that they do not compromise the security of the information entrusted to them. Host components are interconnected via a multilevel communications subnet (which may itself be composed of components and simple communications channels. Subjects within one component can (by interacting with the appropriate device drivers) cause information to be exchanged between components in a secure way.

Note that a point-to-point connection may be abstracted as a pair of devices (one at each end) linked by a communication medium. A broadcast channel may be abstracted as a set of devices (one for each host) linked by a shared communication medium. The hypothesized network may contain both single-level and multi-level connections.

B.4.5. Simulation of the Network on the Monolithic System

The proposed system may be simulated in a very natural way on the evaluated monolithic computer system.

Each component subject (in the network) is simulated as a single subject (on the monolithic target system.) For reasons that will become clear later, all of the network subjects within a single component are allocated to a single

processor of the monolithic system, and it is assumed that there is a processor available for each network component.

All of the communication devices are provided as I/O devices within the computer system; single- or multi- level as appropriate. For each device, it is supposed that there is a server subject, which correctly implements the protocol ascribed to the communication channel and, for multilevel devices, has the trust properties required to function as a trusted subject. As each device is local to a processing node in the network system, it is made local to the associated processor in the monolithic computer system (i.e., it is accessible only by that processor).

Finally, the I/O devices are linked using the appropriate physical media, (which is considered to be external to the system): in pairs, for point-to-point channels, and in sets, for broadcast channels.

The simulation is now an accurate representation of the hypothesized network. Since it runs on an evaluated monolithic system, it is secure to the degree of assurance ascribed to the monolithic system, subject, of course, to the provision of appropriate levels of communication security to the various communications channels. The Criteria Interpretations provided in the TNI may be viewed (for the higher evaluation Classes) as specifying the characteristics a network must have to be simulated in the way described.

B.4.6. Transformation of the Monolithic Simulation to a ----- Distributed System -----

It is instructive to examine certain of the properties of the network simulation.

It may be observed that there are no application memory segments shared by subjects allocated to different processors. This stems from the allocation of all subjects within a single network component to a single processor of the monolithic system, and from the rule (for the network) that no subjects access objects in a different component.

Furthermore subjects executing on different processors do not utilize any of the inter-process communications mechanisms provided by the TCB; all inter-processor communication is provided by means of the I/O device protocols embedded in the I/O device drivers, which are part of the TCB. Moreover, the (correct) operation of these protocols does not depend upon the sharing of memory since they were usable in the network being simulated, and thus presumably provide for the cooperation of remote devices coupled by a shared physical medium.

Thus, outside of the security kernel, no memory segments are shared by any two processes running on different processors. Assuming that each processor has local memory, all application segments may be moved (without effect) to the appropriate processor-local memory address space. Supposing the TCB code is ``pure'' (i.e., re-entrant), complete copies of the TCB code may also be removed to the local memory address space of each processor without effect. Similarly, internal TCB data structures that have elements that are accessed only by a single processor can be removed to the local memory of that processor without effect.

It may be noted (based upon available worked examples) that the only data structures within the TCB, that must be shared by processors, are those representing resources shared by processes running on the various processors. However, in the simulation just described, there are no such shared resources. Devices in the network are local to network components and are therefore accessed only by subjects running on one processor in the computer system. No inter-process communication takes place between processors (it is all via external communication channels), and the only shared global memory required is for the table controlling global memory allocated to the subjects, of which there is none.

Thus, in the particular network simulation described, despite the potential for shared resources assumed by the underlying TCB, that potential is never exercised. The partitioning of code and data described allows the internal restructuring of the TCB in such a way that the TCB is partitioned and removed to processor local memory throughout, with no residual code or data in global memory. This internal restructuring in no way affects the operation of the system and in no way impacts its compliance with the TCSEC Criteria (for the specific application).

Another result of the described partitioning and localization of the TCB is that no communication ever takes place over the system bus: all of the TCB tables may be locked locally so that no inter-processor communication within the TCB is required, and there are no global memory segments. It follows that the bus may be completely severed without affecting either the operation of the system or its compliance (in this particular case) with the Criteria. An interesting observation is that no single step in the restructuring described can be regarded as changing the fact that the collection of processors is utilizing a single TCB, which is compliant with the Criteria. It is this observation that impels one to conclude that a single TCB can be properly implemented as a collection of TCB partitions.

The resulting partitioned TCB is now examined. Within the TCB are a set of (trusted or untrusted, as appropriate) I/O device drivers, one for each I/O device. As constructed, a particular device is utilized only by the subjects being executed by a single processor: the driver subject for that device exists, but is quiescent, on all other processors (because none of the application subjects are attempting to utilize that device). The driver subject, its code, and its data may therefore be removed from the TCB partitions for all of the processors except that for which the device is local. Again, the system remains a valid interpretation of the model, and remains compliant with the Criteria.

The resulting system still has only one TCB, partitioned among a number of asynchronous processors, with the code and data for supporting various devices provided only within those TCB partitions where they are needed to support local devices. The only links between the physical processors are the various single- and multi-level communication channels provided. These channels are afforded, it has been assumed, the appropriate levels of physical security by communications security techniques, just as they would be if they were media connecting a computer to a remote terminal: the provision of this physical security is an axiom in the context of evaluating the validity of the system from a ``computer security'' point of view. (This is discussed fully in section B.6, as the importance of communications security techniques in the context of a network of systems must not be trivialized).

Each processor, and its associated devices, is now packaged in a separate physical box. There is now little difference between the hypothesized ``monolithic computer system'' (with an admittedly very specialized application running on it) and the network originally hypothesized. The single TCB of the partitioned ``monolithic system'' remains a single TCB, which has, however, been transformed into a collection of TCB partitions, each of which is responsible for enforcing access control policy within its ``local partition'', or component.

The TCB in a particular box may now be replaced by an equivalent TCB (that is, a TCB with the same top-level specifications, and with an equivalent degree of assurance) without impacting the overall security of the system or its adherence to the TCSEC Criteria. In fact, both the hardware and software TCB bases within a partition could be replaced, as long as the replacement has the same (or greater) evaluation class and completely honors the interface protocols (and thus, for example, correctly receives and transmits

labeled datagrams) defined for the devices connecting it with the other processors.

Finally, the particular Formal Security Policy Models upon which the TCBS within each box are based might be allowed to differ without adverse impact, so long as each model used was a valid representation of the single Network Security Policy to be enforced, as allocated to the activities of the application subjects within the box.

B.4.7. Conclusions Regarding the Simulation Argument

This informal argument shows how a network of processing nodes, which are ``locally autonomous'' (with respect to their enforcement of a global Security Policy for access controls), can be simulated upon a clearly evaluable monolithic system with a security kernel, and, in turn, how that system can be physically partitioned into a confederation of components, each with its own TCB partition. The resulting system is the originally desired network in all of its essential features, and is clearly in harmony with the intent of the TCSEC Criteria. This argument provides an intuitive basis for the interpretation of the Criteria provided in the TNI. It also shows the sense in which the collection of NTCB partitions may be viewed as forming a single

NTCB: there is a single NTCB because there is a single Security Policy for the network, which is locally enforced by each NTCB partition upon its local subjects and objects (i.e., upon the resources it controls).

Of significance for the design and evaluation of networks targeted for the higher evaluation classes is the fact that under the assumptions that an overall Network Security Policy has been defined, and that the communication channels between components function correctly, (i.e., maintain the proper associations between labels, user identifications, clearances, and names of objects), there is no compelling reason to insist that the required assurance for each processing node be obtained using the same Formal Security Policy Model.

B.5. Cooperation Among Partitions

In this section we focus on that part of the NTCB outside the security kernel, i.e., that part involved in the implementation of supporting policies and typically carried out by trusted subjects. Some non-kernel NTCB functions are essentially the same as those normally provided in a non-networked trusted computer system, such as login authentication of local users. Such functions in an NTCB partition can be understood in terms of the services they perform

within the network component.

Other non-kernel NTCB functions provide distinctively network-related services that can best be understood in the context of the network security architecture. We shall refer to these as trusted network services. Very often, an essential task of these functions is to implement a protocol for conveying security-critical information between trusted subjects in different components. The trusted protocol is not an end in itself, but a means to accomplish services for which cooperation between NTCB partitions is needed. A simple example would be the need to change the security level of a single-level communications channel. While each component could internally relabel the I/O device connected to its own end of a channel, a trusted protocol is required to coordinate the changes.

In this section there will be two brief examples illustrating the relationship between a network security architecture and an associated trusted network service. One example network uses trusted network interface units and protected wire distribution, and the other uses end-to-end encryption. After the examples, design specification and verification of trusted network services will be discussed.

B.5.1. Trusted Interface Unit Example

Consider a network in which untrusted hosts operating at various single security levels communicate through trusted network interface units (TIU's) that send and receive labeled messages in the clear over a protected communication subnet. The function of a TIU is to place message sensitivity labels on outgoing messages, and to check labels on incoming messages, so that hosts may send and receive only messages labeled in accordance with their accreditation.

Because the communication subnet carries messages at all levels, the I/O device connecting any TIU and the subnet is single-level system-high. But the connection between any TIU and its host is at the level of the host. Thus, a TIU for a low-level host must contain a trusted subject that reads high and writes low.

There is a trusted protocol in this example, though it is relatively trivial, since it merely identifies a header field in each message that should contain a sensitivity label, and perhaps also a checksum to guard against transmission errors. A protocol of this kind is required whenever information is exported or imported over a multilevel communications channel. See section 3.1.1.3.2.1 of Part I of this document.

B.5.2. End-to-End Encryption Example

Consider a network in which hosts operating at various security levels communicate through trusted front-end processors (TFE's) that send and receive encrypted messages over a public communication subnet. Suppose that the TFE's obtain encryption keys at the level of the information to be protected from a central key distribution center (KDC) supporting the various security levels of the network, attached to the network in the same way as a host. A key is sent from the KDC to a TFE upon request, using an appropriately certified protocol that authenticates both the requester and the new key.

The purpose of key distribution is really to support a trusted local service within the TFE, namely, the ability to transform classified messages from the host into unclassified encrypted messages suitable for transmission over the subnet. In other words, there is a trusted subject that reads high and writes low.

Part of the trusted network service is implemented within the KDC, which must generate new keys for the level of information being communicated, and must also decide, on the basis of an access control policy, which TFE's may share keys. A single level subject in the KDC at the level of the information which the key is for does not necessarily require privileged treatment from the kernel in the KDC; however, such trusted network service subjects at the various levels must correctly implement a certain policy and a certain protocol.

B.5.3. Design Specification and Documentation

To obtain the level of assurance needed for systems of Evaluation Class A1, a formal top-level specification (FTLS) of the NTCB is required, including a component FTLS for each NTCB partition. As in the case of stand-alone computer systems, non-kernel portions of the NTCB must be specified even though they support policies that are not part of the access control policy represented in the formal security policy model. In particular, software supporting trusted network services in each component must be specified.

Where a trusted network service supporting the mandatory policy depends on a protocol, the protocol will necessarily appear in FTLS of some component(s) of the network. As a minimum, the role of the trusted subject in each NTCB partition will be implicit in each component FTLS. Trying to understand a protocol by looking at each participating process separately, however, is like trying to read a play in which the lines have been sorted by character. For pur-

poses of documentation, it is desirable to provide a representation of each trusted protocol in a fashion that exhibits the interactions between participants, and the correspondence between this representation and the relevant parts of component FTLS's should be shown.

Just as the FTLS of a stand-alone TCB contains representations of operating system conceptual entities, such as processes, devices, memory segments, and access tables, the FTLS of an NTCB contains representations of protocol entities and concepts, such as connections, where they occur, such as in trusted network service specifications.

In the end-to-end encryption example, correspondence of the FTLS to the trusted network services supporting policy should include the demonstration of at least the properties that all data transmitted over the communication subnet is encrypted with the proper key (e.g., for the correct security level), and that the KDC allows keys to be shared only in accordance with its access control restrictions. Both properties might be stated and proved in terms of connections between hosts. In the trusted interface unit example, the correspondence should show that each TIU marks and checks message labels in accordance with a given host label.

B.5.4. Summary

Some non-kernel NTCB functions in a network may be characterized as trusted network services. They provide trusted protocols to implement security-critical cooperation between trusted subjects in different NTCB partitions. Showing correspondence between the FTLS for these services and their supporting policies implies proving certain properties, expressed in terms of network-specific concepts, which convey essential features of the network security architecture.

B.6. Communication Channels Between Components

In this section the communication channels used to connect components are examined more closely, with the goal of understanding when the characteristics of a particular channel are relevant to the security characteristics of the system, how the characteristics of such a channel are to be evaluated and related to the overall evaluation of the network, and those factors that must be deferred to the assessment of the adequacy of the network to support a particular application of it preceding its accreditation.

The discussion is organized into the following major parts. In section B.6.1, the notion of a ``communication channel'' is related to the technical terminology provided

by the TCSEC Glossary. In section B.6.2, the notion of a ``Security-Compliant communication channel'' is defined. The remaining parts of the section discuss the important cases of channels that are single-level and multilevel (in the mandatory policy sense).

B.6.1. Basic Notion of A Communication Channel

For the purposes of the TNI the network is viewed as a system of components, connected (at the physical layer) by communication channels. The term ``communication channel'' is used as a refinement of the term ``channel'', defined in the TCSEC Glossary as ``an information transfer path within a system.'' The term may also refer to the mechanism by which the path is effected. It is further required, for the purposes of applying the TNI to a network, that the network architecture be formulated in sufficient detail that all communication channels are Security-Compliant as defined below.

``Point-to-point'' communication channels are discussed first. The notions of ``communication channel'' and ``I/O device'' are distinct: a point-to-point communication channel is viewed as consisting of two I/O devices (each local to the component it is attached to) coupled by a communications medium (which may in reality consist of a complex arrangement of internal devices, switches, and communications links). From the point of view of the components, information is transmitted via the transmitting and receiving devices in a sufficiently error-free, physically secure fashion to merit the particular labels associated with the device. It is, of course, the concern of both the sponsor and evaluator of a particular network to confirm that this condition is met to an appropriate level of assurance, depending upon the security policy allocated to the channel. This requirement, which is a boundary condition upon which the evaluation of the NTCB partition itself, will typically be met by a combination of error-detection and recovery techniques, cryptographic techniques, and other communications security techniques as addressed in Section 9 of Part II.

(Note: Figure not included)

Figure B1. Point-to-point communication channel.

For example, two processing nodes connected by a single channel would be modeled as shown in Figure B1. Here, we would speak of processing component P1 using I/O Device D1 to communicate, via I/O Device D2, with processing component P2. D1 and D2 are assumed to be linked by some sort of physical medium M (perhaps a set of wires, perhaps something

more complicated.) Subject S1 in component P1 may transmit information to a subject S2 in P2 as follows: each subject obtains an object of the appropriate class for use as a buffer. Each subject attaches its locally available device. Subject S1 in P1 then transmits the information in its buffer to D1; subject S2 receives the information via D2 in its buffer. Note that in this description, it was quite unnecessary to introduce the notion of either a shared object or a shared device. Of course, the details of the inter-communication will depend upon a shared communications protocol.

Broadcast communication channels are only slightly different, from the point of view adopted within the TNI, from point-to-point communication channels. Instead of a pair of I/O devices linked by a physical medium, there is a set of I/O devices linked by a physical medium. Each device may be a receiver, a transmitter, or a transceiver in nature. It is assumed that anything transmitted by a transmitter can be received by any receiver. (It is, of course, determined by the communication protocols being executed by the various devices whether reception actually results in any meaningful action by a particular receiver.)

B.6.2. Security-Compliant Channels as the Basis for Evaluation

Communication channels in trusted network architecture must be Security-Compliant. A channel is Security-Compliant if the enforcement of the network policy depends only upon characteristics of the channel either (1) included in the evaluation, or (2) assumed as an installation constraint and clearly documented in the Trusted Facility Manual. The first approach tends to produce evaluated network systems whose security characteristics are relatively immune to installation or configuration choices. The second approach yields evaluated network systems whose security is more strongly conditioned upon the appropriateness of installation or configuration choices; however, the conditions and limitations of the evaluation are clearly documented.

The overall security of the network can be assessed by verifying the correctness of the NTCB partitions (an evaluation issue) and by verifying that the required environmental constraints documented for all communications channels are, in fact, met by the installation (an accreditation issue). The thrust of this section is to show that channels that are not Security-Compliant may be reduced to Security-Compliant channels so that the resulting architecture will support a viable network evaluation. Three general techniques are available for rendering a channel Security-Compliant: 1) the

utilization of the channel for security-critical transmissions may be restricted by using controls internal to the NTCB partitions of the components linked by the channel; 2) end-to-end communications technologies (such as encryption) may be installed and evaluated as part of the linked NTCB partitions to eliminate the influence of the channel's physical environment on the security properties of the channel; and 3) constraints on the intrinsic characteristics assumed for the channel may be documented in the Trusted Facilities Manual. The last approach, in effect, reserves determination of the adequacy of a particular channel to the accreditor: the evaluation proper will be based upon a communications channel, which will be assumed to have the desired characteristics.

The evaluation effort is focused upon establishing the correctness of the technique, or combination of techniques employed. The adequacy of the mechanisms is an accreditation issue. For example, the issues related to the adequacy of data confidentiality service are discussed in Part II.

A channel can be made Security-Compliant by using a combination of the above techniques: cryptographic sealing, for example, addresses the issues of both prevention of unauthorized modification and error-detection. In evaluating each channel, three vulnerabilities related to external environmental factors and one related to internal exploitation must be addressed. They are as follows:

1. Communication security - unauthorized disclosure or modification of sensitive information in transit
2. Communication reliability - unreliable delivery of information, (e.g., non-delivery, misdelivery, and delivery of erroneous data) the delivery of which is required for the correct operation of the NTCB (such as audit records or inter-partition security coordination)
3. Communication fidelity - changes to security-critical data, such as transmitted security labels, due to noise. (Note that changes due to unauthorized modification are categorized as a communications security problem)
4. Covert signaling - manipulation of the channel mechanisms to signal information covertly

The use of a channel as a covert signaling mechanism will be evaluated in the normal course of events (if required by the Criteria) when the required covert channel analysis of the channel drivers, which are part of the linked NTCB partitions, is performed. See the Covert Channel Analysis section in Part I. Techniques for addressing

the remaining three vulnerabilities are listed below.

The first vulnerability, to the security of sensitive information in transit, must be addressed by one or more of the following techniques:

1. Documenting a constraint in the Trusted Facilities Manual that the installed channel be completely contained within an adequate security perimeter (thereby deferring an assessment of compliance to accreditation)
2. Providing, for the channel, suitable end-to-end communications security techniques which are documented and evaluated as part of the NTCB partitions linked by the channel
3. Restricting utilization of the channel to the transmission of non-sensitive information by means of controls internal to the NTCB partitions linked by the channel

Vulnerability of a channel to the unreliable delivery of security-critical information must be addressed by one or more of the following techniques:

1. Documenting a constraint in the Trusted Facilities Manual that the channel be comprised of intrinsically reliable media and devices (thereby deferring an assessment of compliance to accreditation)
2. Providing for the channel suitable end-to-end protocols for the reliable transmission of information within the NTCB partitions coupled by the channel, which will thereby be evaluated for correctness
3. Restricting use of the channel to transmissions, the delivery of which is not critical to the functioning of the NTCB, by means of controls internal to the NTCB partitions linked by the channel, which will thereby be evaluated for correctness

Vulnerability of a channel to noise, which may compromise the correctness of security-relevant data (such as security labels) must be addressed by one or more of the following techniques:

1. Documenting a constraint in the Trusted Facilities Manual that the channel be comprised of intrinsically noise-free media and devices (thereby deferring an assessment of compliance to accreditation)
2. Providing for the channel suitable end-to-end noise reduction techniques within the NTCB partitions

linked by the channel, which will thereby be evaluated for correctness

3. Restricting use of the channel to transmissions, the noise-free delivery of which is not critical to the functioning of the NTCB, by means of controls internal to the NTCB partitions linked by the channel, which will thereby be evaluated for correctness

Three example scenarios are provided below, showing how these techniques might be employed.

Example A. Two loosely-coupled trusted coprocessors, one in active use and the other in ``hot standby'', are to be linked by a dedicated communications channel. Significant amounts of dynamic, security-relevant data will be exchanged over this channel. The channel must be trusted to preserve label integrity and provide reliable and noise-free delivery of security-critical data. Noise is not a design issue. The channel must reside in a physically secure environment.

The simplest evaluation strategy would be to document the required environmental constraints in the Trusted Facilities Manual: that the channel be placed within the ``system-high'' security perimeter, and that it be comprised of intrinsically reliable and noise-free media and devices. During evaluation the proper documentation of these constraints would be verified. Compliance of the selected channel in the physical installation to them would be an accreditation issue which, in this case, would (apparently) be easy to verify. This evaluation approach would have the advantage of allowing replacement of the original channel with a higher-performance version without inducing re-evaluation (although the system would have to be accredited again).

Example B. Numerous single-level hosts (at several levels) are interconnected via a multilevel packet switch, which emulates multiple point-to-point networks between communities of hosts of the same level. Communications between host and packet switch are in the clear the sensitivity level of each host is determined at the switch by internally labeling the hard-wired communication ports. The communication channel must be secure (separation of data of different levels must be maintained), but it need be neither reliable nor noise-free (from the point of view of security).

Two quite different approaches might be regarded as suitable for the evaluation of this architecture. In the first, (and most natural), the architecture would be reformulated to show the packet switch as a network component,

connected to each host with a single-level channel. Network documentation would then indicate that each of the new channels be constrained to be located within an appropriate security perimeter (so that physical security is maintained), and that no security-critical information requiring either reliability or fidelity is transmitted over them. The second assertion would be verified during evaluation, and the first during accreditation.

A second (radical) approach would be to insist that the packet switch is part of the communication channel. In this case, it is difficult to see how the required Security-Compliance is to be attained while encompassing the packet switch within the evaluation, as end-to-end techniques are not in use, and it is obvious that sensitive information is being transmitted. The sponsor could document a constraint upon the interconnection of hosts that the (nominally point-to-point) channels be each confined within a security perimeter, thereby excluding the packet switch from evaluation, but it is also then dropped from the description of the network being evaluated, and is replaced by ``nominally'' Security-Compliant point-to-point channels, documented in the Trusted Facilities Manual. The decision to use a particular multilevel packet switch to meet the documented requirement for an adequate security perimeter between the point-to-point virtual channels would then be the responsibility of the accreditor of such a system alone. In effect, such a strategy when applied to the system described is a decision to use the techniques described in Appendix C (the system actually evaluated is an uninteresting subset of the originally envisioned system)

Example C. Two trusted multilevel systems are to conduct file transfers over a channel, which is intrinsically noisy, unreliable, and insecure. The data is to be encrypted and cryptographically sealed. Reliable transmission is enforced by non-NTCB software. (This is not security-relevant, however, because the loss of a data packet is not an insecurity).

The communications security and cryptographic sealing techniques must be included within the evaluated NTCB partitions. Assessment of their correctness will be part of the evaluation, and assessment of their adequacy, based upon the true sensitivity of the information transferred, will be part of the accreditation. In order to address channel reliability, the NTCB internal control mechanisms preventing the utilization of the channel for security data needing to be transmitted reliably must be documented and evaluated (in this case, the argument is probably the degenerate case: that no such information exists.) See, for example, the Encryption Mechanism section in Part II.

B.6.3. TCSEC Criteria for Multilevel Communication Channels

In this section, those TCSEC Criteria relevant to the utilization of communication channels within networks are examined, from the point of view of countering internal threats. As the Criteria for Class A1 are the most stringent and are a superset of the requirements for other classes, they are the basis for the discussion.

The Class A1 Criteria (Section 4.1.1.3.4) requires that ``the TCB shall support the assignment of minimum and maximum security levels to all attached physical devices.'' The basis for making this designation is also stated in Section 4.1.1.3.4: ``these sensitivity levels [i.e., for devices] shall be used by the TCB to enforce constraints imposed by the physical environments in which the devices are located''.

In the case of a communication channel connecting components of a network, the ``physical environment'' must be interpreted to include the environment of the devices and the medium linking them. The range of access classes (from the Network Mandatory Access Control Policy) to be assigned to the channel must take into account the physical security afforded to the medium, the communications security techniques that have been applied to the secure information being transmitted through the medium, the physical accessibility of the devices involved in the channel, and the intended use of the channel from an architectural point of view for the network. Within these constraints, the Criteria cited requires that the devices comprising the channel be appropriately labeled (with, it may be inferred, locally ``appropriate'' internal labels). For example, a particular channel may be designed to support the transmission of UNCLASSIFIED through SECRET information. The receivers and transmitters coupling the transmission medium to the hosts (assuming all hosts receive and transmit at all levels) must be labeled within each host with whatever the local internal labels are designating the UNCLASSIFIED through SECRET range. That such labels exist is guaranteed by the requirement to interpret properly a network Security Policy for each NTCB partition.

In addition to the labeling of devices coupling network processing components to the communication channels which may exist, the TCSEC requires, for multi-level channels, that all information exported to, and imported from, the channel be properly labeled: ``when exported by the TCB, sensitivity labels shall accurately and unambiguously represent the internal labels and shall be associated with the information being exported'' (Section 4.1.1.3.2); furthermore, ``when the TCB exports or imports an object [sic] over a multilevel channel, the protocol used on that channel

shall provide the unambiguous pairing between the sensitivity labels and the associated information that is sent or received'' (Section 4.1.1.3.2.1). The interpretation of these Criteria appears to be straightforward: in the context of a network communication channel, they imply that information be properly labeled when it is exported, that there be a shared protocol between the exporting and importing devices which unambiguously and correctly maintains the label-

information association, and that the resulting imported label be honored by the receiving NTCB partition. Note that the requirement for integrity relative to label-data associations is clearly stated and need not be hypothesized as a requirement specific to networks.

B.6.4. Single-Level Communication Channels

The Criteria states that ``the TCB shall support the assignment of minimum and maximum security levels to all attached physical devices'' which ``enforce constraints imposed by the physical environments in which devices are located'' (Section 4.1.1.3.4). Note that this capability is required to exist for all devices, whether ``single-level''

or ``multilevel''. The distinguishing characteristic of ``single-level devices and channels'' is stated in Section 4.1.1.3.2.2, ``single-level I/O devices and channels are not required to maintain the sensitivity labels of the information they process''. Thus, a device and/or channel which does not support the transmission of labeled information is, by definition, ``single-level''.

There are two cases: the minimum and maximum security levels of the devices coupling the channel to the processing nodes may be all the same, or not.

The case in which all of the minimum and maximum security levels are identical is the normal case (for network communication channels) of a channel which is to transmit information of a single, invariant, sensitivity level.

It is also possible that the minimum and maximum ranges of the various devices associated with a single-level channel are not all the same. In this case, the channel may carry unlabeled information, but of only one sensitivity level at a time. It is the responsibility of each NTCB partition coupled to the channel to prevent the transmission of information of a sensitivity level different from the current level of the channel in accordance with Section 4.1.1.3.2.2, ``the TCB shall include a mechanism by which the TCB and an authorized user reliably communicate to designate the single security level of information imported or exported via single-level communication channels or I/O

devices''. In the context of a partitioned NTCB, this means that single-level channels may be defined as part of the network architecture which can be manually shifted from the transmission of one level of unlabeled information to another by an authorized user. The natural interpretation of the criterion cited above is that a reliable protocol must exist for informing each NTCB partition involved in controlling access to the channel that a change in level has been ordered, prior to the transmission of any information over the channel (so that the ``implied'' label can be correctly assigned by each NTCB partition to information received).

B.7. Miscellaneous Considerations

B.7.1. Reference Monitor, Security Kernel, and Trusted Computing Base

The notion of a ``reference monitor'' is the primary abstraction allowing an orderly evaluation of a stand-alone computer system with respect to its abilities to enforce both mandatory and discretionary access controls for the higher evaluation Classes.

The TCSEC Glossary defines the ``Reference Monitor Concept'' as ``an access control concept that refers to an abstract machine that mediates all accesses to objects by subjects''. Although the reference monitor abstraction includes the notion of protection, the abstraction itself is independent of any particular access control policy. The abstraction assumes that a system is comprised of a set of active entities called ``subjects'' and a set of passive entities called ``objects''. The control over the relationships between subjects and objects, i.e., the access of objects by subjects, is mediated by the reference monitor in such a way that only accesses permitted by the access control policy being enforced, are permitted by the reference monitor. The reference monitor is thus the manager of the physical resources of a system. A distinguishing feature of the monitor is that there is a well-defined interface, or ``perimeter'', between the reference monitor itself and the subjects and objects it controls. To be effective in providing protection, the implementation of a reference monitor must be (1) tamper-proof, (2) always invoked, and (3) simple enough to support the analysis leading to a high degree of assurance that it is correct.

The hardware and software components of a computer system implementing a reference monitor meeting these principles is called a ``security kernel'', defined in the TCSEC Glossary as ``the hardware, firmware, and software elements

of a Trusted Computing Base that implement the reference monitor concept. It must mediate all accesses, be protected from modification, and be verifiable as correct''.

From this definition, it is apparent that a ``security kernel'' (if there is one) is always part of the TCB of a computer system, defined as ``the totality of protection mechanisms within a computer system - including hardware, firmware, and software - the combination of which is responsible for enforcing a security policy ... the ability of a Trusted Computing Base to correctly enforce a security policy depends solely on the mechanisms within the TCB and on the correct input by system administrative personnel of parameters (e.g., a users' clearance) related to the security policy.'' In particular, the TCB includes those mechanisms involved in the implementation of supporting policies, while the (included) security kernel (if there is one) is involved only with the enforcement of access control policies.

B.7.2. Network Trusted Computer Base and Reference Monitor

The notions of a TCB, and, for the higher evaluation classes, security kernel and reference monitor can, with suitable interpretation, be applied directly to the evaluation of trusted networks without significant change. In particular, the Network Trusted Computing Base can be defined as the totality of protective mechanisms within the network (including mechanisms provided by connected host systems) responsible for enforcing the (overall) security policy. Implied in this definition is the strong notion that an evaluated network has a single NTCB, as the NTCB is, by definition, the totality of enforcement mechanisms for the stated policy.

For the higher evaluation classes the TCSEC requires, in effect, that a reference monitor be implemented as part of the TCB. This, at least in theory, is not the only conceivable technology for implementing a highly-assured system (one could envision a completely verified system, for instance); however, it appears to be the only current technology with a proven track record, and within the current state-of-the-art to be able to be implemented economically.

For these reasons, the Part I of the TNI takes a pragmatic stance with regard to specifying interpretations for Trusted Networks: that the TCSEC notions of ``reference monitor'' and ``security kernel'' be applied in the network system context with as little change as possible for the higher evaluation classes. We therefore assume that the

NTCB of a Trusted Network of class B2 or above must contain the physical realization of a reference monitor which mediates all references within the networked system of subjects to objects, is tamperproof, and is small enough, in aggregate, to validate.

B.7.3. NTCB Partitions

The view taken of a ``network system'' throughout the TNI is that the network can be partitioned into ``components'', each of which has various processing and communication capabilities. Given such a decomposition, the functions of the NTCB must be allocated in some coherent way to the various components of the network.

The following terminology is introduced: the totality of hardware, firmware, and software mechanisms within a single network component which is responsible for enforcing the security policy of the network, is called the NTCB partition within that component. As the collection of network components and communication channels is meant to be exhaustive (i.e., every part of the network system, including hosts, is accounted for) and disjoint (i.e., no parts are shared between components), the NTCB partitions collectively are a true partition of the NTCB; they are non-overlapping and complete.

For Mandatory Access Control Policy, a large and useful class of networks can be envisioned, which allows a clean decomposition of the NTCB into NTCB partitions. The resulting partitions can be evaluated relative to enforcement of mandatory access controls using conservative interpretations of the TCSEC, and the correctness of the composition of these components into a network enforcing the overall policy for mandatory access controls is easily confirmable as well. For sponsors wishing to obtain an overall evaluation class for a network system, such a ``partitionable'' network architecture should be chosen.

Concisely, the network architecture must have the following salient features: (1) subjects and objects within multilevel processing components are given the usual TCSEC interpretation; (2) subjects and objects are confined within their individual components, i.e., no subjects may directly access information within a different component

(In practice, this means there are no directly accessible, shared memory registers); and (3) information representing the access-relevant security state of the subjects and objects within a component, is maintained locally by the NTCB partition of that component. (It may be the case that information representing the components state may be distributed to other components for the purpose of overall network

control, recovery, etc. but the decision to permit or prohibit access is always made locally, based on locally available state information.

A network of host processors and peripherals, interconnected according to these rules, may be roughly described, with regard to security, as a network of ``loosely-coupled'' security kernels. Each security kernel is autonomous with regard to the accesses made locally (i.e., within the component whose resources the reference monitor controls). A subject within one component may transmit data, under the control of the two security kernels, to a subject in another component. However, the basic principle to be seen at this point is the following: because all accesses are local (i.e., constrained to be by a subject within a component to an object within that component), all accesses are mediated by the security kernel within a component. Thus, the totality of all of the security kernels within the system is adequate to mediate all accesses made within the system.

B.8. Summary and Conclusions

In this Appendix, the rationale for the partitioning of the NTCB into a set of cooperating, loosely-coupled NTCB partitions has been presented. Each NTCB partition may be viewed as a locally autonomous reference monitor, enforcing the access of local subjects to local objects and devices. Because the partitioning is carefully constrained to reflect the lack of sharing of objects among components, (while allowing the transmission of information between components through shared physical media), the aggregate of locally autonomous NTCB partitions is adequate to mediate all accesses of subjects to objects and thus is adequate to form the basis for a reference monitor for the entire network system. Under the conditions of loose coupling assumed, the specification of a network-wide Security Policy, properly interpreted as an individual Formal Security Policy Model for each component enforcing access controls, suffices to provide the basis for demonstrations that each component meets the requirements of the Security Policy. The postulated network architecture and design (that are presented by the network sponsor) suffices to allow evaluation of the supporting policy capabilities required to attain the targeted evaluation class.

APPENDIX C

Interconnection of Accredited AIS

C.1. Purpose

As was discussed in the Introduction to this document, there are many "networks" that can not be meaningfully evaluated as a "single trusted system" because they are sufficiently complex and heterogeneous that no single evaluation rating can adequately reflect the trust that can be placed in the "network". The purpose of this Appendix is to provide guidance concerning how to interconnect systems in such a way that mandatory security policies are not violated.

C.1.1. Problem Statement

The interconnected accredited Automated Information System (AIS) view is an operational perspective which recognizes that parts of the network may be independently created, managed, and accredited. Interconnected accredited AIS consist of multiple systems (some of which may be trusted) that have been independently assigned accreditation ranges, reflecting the various sensitivity levels of information that may be simultaneously processed on that system. In this view, the individual AIS may be thought of as "devices" with which neighboring systems can send and receive information. Each AIS is accredited to handle sensitive information at a single level or over a range of levels.

An example of when the interconnected accredited AIS view is necessary is a network consisting of two A1 systems and two B2 systems, all of which are interconnected and all of which may be accessed locally by some users. It is easy to see that, if we regard this as a single trusted system, it would be impossible for it to achieve a rating against Part I of this document higher than B2. This might not be an accurate reflection of the trust that could be placed in the two A1 systems and interconnections between them. The single rating of B2 assigned to this network could be misleading.

While it provides much less information about a system than does a meaningful evaluation rating, taking the interconnected accredited AIS view of the network provides guidance on appropriate interconnection strategies.

C.1.2. Component Connection View and Global Network View

There are two aspects of the Interconnected Accredited AIS view of a network that must be addressed: the component connection view and the global network view. These two

views are discussed below and will be examined in greater detail later in this Appendix.

Any AIS that is connected to other AIS must enforce an "Interconnection Rule" that limits the sensitivity levels of information that it may send or receive. Using the component connection view, each component responsible for maintaining the separation of multiple levels of information must decide locally whether or not information can be sent or received. This view, then, does not require a component to know the accreditation ranges of all other components on the network; only of its immediate neighbors (i.e., those with which it can communicate without an intermediary).

In addition to the Interconnection Rule, there may be other constraints placed on a network to combat potential security problems. The global network view is a way of addressing some of the other constraints placed on a network. This view requires one to have a knowledge of the accreditation ranges of all components of the system. These accreditation ranges are taken into account when determining whether or not a component should be allowed to connect to the system. In this way, the potential damage that can occur when information is compromised or modified can be limited to an acceptable level.

An example of a problem for which constraints may be placed on the network is what is called the "Cascading Problem." This occurs when AIS are interconnected in such a way that the potential damage from unauthorized disclosure or modification is above an acceptable level. The network sponsor may wish to limit the damage that can occur by limiting the accreditation ranges of AISs that can be interconnected.

C.2. Accreditation Ranges and the Interconnection Rule

C.2.1. Accreditation Ranges

The AIS accreditation range reflects the judgement of the accreditor on the ability of the component to appropriately segregate and manage information with respect to its network connections in accord with the designated sensitivity levels. An ADP system that has been accredited for (stand-alone) system high operation would be assigned an accreditation range having a single sensitivity level equal to the system high sensitivity level of the system. Such a system is not relied upon to segregate the several actual levels of information processed. All the information exported from such a system must be labeled with the system-high sensitivity label until there is a competent manual review to assign a lower level. A multilevel

(stand-alone) AIS might be assigned an accreditation range equal to the entire set of levels processed. In this case, the label of the exported data is equal to the actual level of the data processed within the accredited range.

In a network context, the accreditation range bounds the sensitivity levels of information that may be sent (exported) to or received (imported) from other components. If a network has only dedicated and system-high components, each component will be assigned single-valued accreditation ranges (i.e., an accreditation range consisting of one sensitivity level).

Consider an example, illustrated in Figure C1, which uses accreditation ranges along with an approach based on the Environmental Guidelines.- Component A is a class B2

system and has an accreditation range of CONFIDENTIAL through SECRET. Component B is a class A1 system and has an accreditation range of CONFIDENTIAL through TOP SECRET. Thus, if Component A has a direct connection to Component B, the accreditation ranges provide a basis for both components to be assured that any data sent or received will not "exceed" (that is, will be dominated by) SECRET in its classification.

Figure C1. Accreditation Ranges Illustrated

(Note: Figure not included)

C.2.2. Interconnection Rule

A multilevel network is one in which some users do not have the necessary clearances for all information processed. A multilevel network therefore is one that processes a range of sensitive information, which must be protected from unauthorized disclosure or modification.

Each component of the network must be separately accredited to operate in an approved security mode of operation and for a specific accreditation range. The component is accredited to participate in the network at those levels and only those levels.

According to this definition, a multilevel network may comprise a mixture of dedicated, system high, compartmented, controlled, and multilevel components, where two or more differ in their classification categories and/or compartments, and/or some users do not have all formal access approvals.

The following requirement must be met in multilevel networks.

- Security Requirements: Guidance for Applying the Department of Defense Trusted Computer System Evaluation Criteria in Specific Environments, CSC-STD-003-85

C.2.2.1. Information Transfer Restrictions

Each I/O device used by an AIS to communicate with other AIS must have a device range associated with it. The device range may be a single level, or it may be a set of levels (in which case the device is referred to as multilevel), and it must be included within the AIS accreditation range.

Information exported or imported using a single-level device is labeled implicitly by the security level of the device. Information exported from one multilevel device and imported at another must be labeled through an agreed-upon protocol, unless it is labeled implicitly by using a communication link that always carries a single level.

Information exported at a given security level can be sent only to an importing device whose device range contains that level or a higher level. If the importing device range does not contain the given level, the information is relabeled upon reception at a higher level within the importing device range. Relabeling should not occur otherwise.

C.2.2.2. Discussion

The purpose of device labels is to reflect and constrain the security levels of information authorized for the physical environment in which the devices are located.

The information transfer restrictions permit one-way communication (i.e., no acknowledgements) from one device to another whose ranges have no level in common, as long as each level in the sending device range is dominated by some level in the receiving device range. It is never permitted to send information at a given level to a device whose range does not contain a dominating level.

It is not necessary for an AIS sending information to another AIS through several other AISs to know the accredi-

tation range of the destination system, but it may be beneficial to network performance; if the originator knows that the information cannot be delivered, then it will not try to send it and network resources will not be used fruitlessly.

In the case of interconnected accredited AISs, the accreditation of a component system and the device ranges of its network interface devices are set by a component administrator in agreement with the network administrator. These ranges are generally static, and any change in them is considered to be a reconfiguration of the network.

In summary, then, if the Interconnection Rule is followed, information will never be improperly sent to a component that is not accredited to handle information of that sensitivity.

C.3. The Global Network View

The above rule applies for communication between any two (or more) accredited systems. However, it does not enforce any of the additional constraints that may be placed on a network. Even when all components have been evaluated (either against the TCSEC, or against Appendix A of this document), and the interconnection rule is followed, there may be other potential security problems. In order to address these problems, it is necessary to adopt a global view of the network. That is, it is no longer determinable locally whether or not a constraint is being satisfied.

Two global concerns will be discussed below. One concern is the propagation of local risk; the other is the cascading problem.

C.3.1. Propagation of Local Risk

The Environmental Guidelines recommend minimum classes of trusted systems for specific environments. The recommendations represent an informed technical judgment on the minimum architectural requirements and assurance appropriate to counter a specific level of risk.

In many cases, operational needs have led to the accreditation of systems for multilevel operation that would not meet the requirements of the recommended class. While the increased risk may be accepted by the users of a particular AIS, connection of such an AIS to a network exposes users of all other AISs in the network to the additional risk.

Consequently, when an unevaluated AIS, or one that does

not meet the class recommended for its accreditation, is proposed for connection to a network, constraints should be considered, such as one-way connections, manual review of transmissions, cryptographic isolation, or other measures to limit the risk it introduces.

C.3.2. The Cascading Problem

One of the problems that the interconnection rule does not address is the cascading problem. The cascading problem exists when a penetrator can take advantage of network connections to compromise information across a range of security levels that is greater than the accreditation range of any of the component systems he must defeat to do so. Cascading is possible in any connected network that processes a greater range of security levels than any one of its component systems is accredited to handle, and it is possible in others as well.

As an example of the cascading problem, consider two systems, each of which is accredited to handle two adjacent classifications of information, as shown in Figure C2. System A processes SECRET and TOP SECRET information, and all users are cleared to at least the SECRET level. System B processes CONFIDENTIAL and SECRET information, and all users are cleared to at least the CONFIDENTIAL level.

While the risk of compromise in each of these systems is small enough to justify their use with two levels of information, the system as a whole has three levels of information, increasing the potential harm that could be caused by a compromise. When they are connected so that SECRET information can pass from one to the other, a penetrator able to defeat the protection mechanisms in these systems can make TOP SECRET information available at the CONFIDENTIAL level.

Figure C2. Cascade Problem, Illustration 1

(Note: Figure not included)

Consider this chain of events: a penetrator (1) overcomes the protection mechanism in System A to downgrade some TOP SECRET information to SECRET; (2) causes this information to be sent over the network to System B; and (3) overcomes the protection mechanism in System B to downgrade that same information to CONFIDENTIAL. This is the cascading problem.

C.3.2.1. Problem Identification

There are various approaches, with different degrees of complexity and precision, for recognizing a potential cascading problem. Two of these approaches will be addressed in this Appendix. The first is a fairly simple test that can ensure that a network does not have a cascading problem:

the nesting condition. The second, discussed in Section C.4, is a less conservative but much more complex heuristic that takes into account the connectivity of the network and the evaluation classes of the component AIS.

The nesting condition is satisfied if the accreditation ranges of every two AISs are either disjoint (have no level in common) or nested, i.e., one is included within the other. In most cases, the nesting condition is enough to determine whether there is a cascading problem. However, this is a somewhat conservative test; there are cases where the nesting condition fails, but there is actually no cascading problem.

Example 1: Consider the situation illustrated in Figure C1. The accreditation range of Component A is nested within that of Component B (i.e., C-S is completely contained within C-TS). Therefore, the nesting condition is satisfied, and there is no cascading problem.

Example 2: Consider the situation illustrated in Figure C2. The accreditation ranges of System A and System B are not disjoint; neither is one completely contained within the other. Therefore, the nesting condition fails, and a cascading condition is indicated.

Example 3: Consider the situation illustrated in Figure C3. Again, the nesting condition does not hold, because the accreditation range of System B is neither disjoint from nor contained in that of Systems A and C. A cascading condition is thus indicated. However, it will be shown below in this Appendix that Figure C3 actually does not contain a cascading condition, due to the presence of the end-to-end encryption devices.

C.3.2.2. Solutions

When a cascading problem is to be addressed, there are several ways to do so. One solution is to use a more trusted system at appropriate nodes in the network, so that a penetrator will be forced to overcome a protection mechanism commensurate with the seriousness of the potential compromise. In the example depicted in Figure C2, If either

system A or system B is evaluated at class B3, which is sufficient according to the Environmental Guidelines for a range of TOP SECRET to CONFIDENTIAL, then the penetrator is presented with an acceptable level of difficulty.

Another possible solution is to eliminate certain network connections, either physically or by means of end-to-end encryption. End-to-end encryption allows hosts that need to communicate to do so, while eliminating additional unnecessary cascading risk on the path from one to the other.

Figure C3. Cascade Problem, Illustration 2 (End-to-End Encryption)

(Note: Figure not included)

In Figure C3, suppose that System A needs only to communicate with System C, and B is just an intermediate node. The possible cascade from TOP SECRET in A to CONFIDENTIAL in B can be avoided by applying end-to-end encryption from A to C, since encrypted data from A can be released at the CONFIDENTIAL level in B without compromise.

Note that end-to-end encryption is of no help in the Figure C2 example, since the systems participating in the cascade were required to communicate.

In some situations where the potential for a cascading problem exists, the risk of its occurrence is actually not significant. A penetration making use of network connections, as described above, generally requires coordination between attacks on two connected systems. It may be possible to determine, in individual cases, that opportunities for this kind of coordination, in the form of common software or common users, are unlikely. One might then disregard cascading over the connections in question.

On a more global scale, one might divide the network into communities, with respect to the possibility of cascading. If connections between one community and another were believed not to support a cascade threat, then a cascading analysis would be performed only within each community.

C.3.2.3. Networks of Evaluated Systems

If the systems to be interconnected can be assigned evaluation classes, the ratings of these systems can be used as input to analysis procedures for detecting the cascade problem and testing proposed solutions. The first step in developing analysis procedures is to define formally the

conditions necessary for the absence or presence of a cascade problem.

An assertion called the Cascade Condition will be defined below. A network satisfying the Cascade Condition does not have a cascade problem. This condition is stated in terms of the evaluation ratings of the interconnected systems and the direction and sensitivity level of connections between them.

Some definitions are needed in order to state the cascade condition formally. The terminology given below is meant to be used only in the context of this section.

A protection region is a pair (h,s) such that h is a network component and s is a sensitivity level processed by component h .

A step is an ordered pair of protection regions $(h_1,s_1), (h_2,s_2)$ such that either

1. $s_1 = s_2$ and h_1 sends to h_2 at level s_1 (a network link), or
2. $h_1 = h_2$ (an information flow within a component).

A path is a sequence of protection regions such that each consecutive pair of regions is a step.

A path is a sequence of protection regions that may be traversed, step by step, by data. A step along a network link is possible either when there is a direct communications link from one component to the other carrying information at a given level, or when there is an indirect, end-to-end encrypted connection in which intermediate components are unable to read the data. A step between two regions in the same component may be (but does not have to be) a covert channel.

Given a host h , let $L(h)$ be the minimum clearance of users of h . Given a sensitivity level s , one can use the Environmental Guidelines to determine the minimum evaluation class $C(s,h)$ required for a system with the associated risk

index. The requirement for open environments should be used unless all systems on the path are closed. Note that $C(s,h)$ will be at least $B1$ if the risk index associated with s and $L(h)$ is greater than zero.

With these definitions, we can now state the Cascade Condition:

For any path $(h1,s1), \dots, (hn,sn)$ such that $sn = L(hn)$ and $C(s1,hn)$ is at least $B1$, there must exist at least one step $(hi,si), (hi+1,si+1)$ such that $hi = hi+1$, the evaluation class of hi is at least $C(s1,hn)$, and si is not dominated by $si+1$.

This condition can be paraphrased by saying that every path that might compromise data of level $s1$ by making it available to an insufficiently cleared user of hn must overcome the protection mechanism in a component of class at least $C(s1,hn)$.

C.4. EXAMPLE: An Heuristic Procedure for Determining if an Interconnection Should Be Allowed

There should be some way of determining whether a system has a risk index that is too great for its evaluation rating (and the evaluation ratings of its components). Given the goal of not allowing a greater risk than is recommended by the Environmental Guidelines, the following is an heuristic algorithm that has been developed to examine systems and determine if they fall within the bounds prescribed by the Environmental Guidelines. (In formal terms, this algorithm is an approximate test for the Cascade Condition, described above.) It should be noted that this algorithm is not intended to be prescriptive: it is merely one way of examining the problem. There are doubtless many other ways that are just as valid.

Furthermore, as any heuristic algorithm, this cannot be

derived from first principles. It has been derived largely through trial and error; it produces reasonable results (e.g., it disallows systems when it seems prudent; it recommends levels of security that are consistent with the Environmental Guidelines).

This algorithm should not be taken to be anything more than intended; it does not magically solve all network security problems. It does, however, provide useful guidance and recommendations as to the prudence of interconnecting various systems.

The following describes an algorithm for determining whether or not a given network, composed of evaluated components, meets the risk categories of the Environmental Guidelines. The algorithm is based on the idea of dividing up a network into groups (where a group is defined to be a group of components that can potentially exchange information i.e., send and receive data at a common sensitivity level, and have an evaluation Class at or below a given level).

The risk presented by any given group can be compared to the maximum allowed risk, as defined by the Yellow Book for a system at the given evaluation class, to determine if any community presents an unacceptable risk.

1. Create a Network Table listing all components within the network. This table, illustrated in Table C1, should include for each component the following information: Component ID, Evaluation Class, Range of Security Classifications at which the component sends data to the network, List of Security Classifications at which the component receives data from the network, Maximum of (highest level of data received from network and highest level of data processed by component), Minimum of (clearance of the user with the lowest clearance of the users with direct access to the component and lowest level of data sent to the network from the component).

Table C1. Example Entry:

(Note: Table not included)

2. Produce a Network Table Evaluation Class, a Network Table Maximum and a Network Table Minimum. The Network Table Evaluation Class will be the highest evaluation class of any component listed in the table. (In Table C1, this is A1.) The Network

Table Maximum will be the maximum of the Maximums associated with all the components listed in the table which send data to the network. (This is determined by taking the highest entry in the "Maximum" column; in Table C1, it is TS.) The Network Table Minimum will be the minimum of the Minimums associated with all the components listed in the table which receive data from the network. (This is determined by taking the lowest entry in the "Minimum" column; in Table C1, this is FOUO.)

3. If the Network Table Evaluation Class is greater than B1, (i.e., A1, B3, or B2) then tables for each evaluation class lower than the Class of the Network Table, must be produced, down to table(s) for the C1 class. These tables will be produced for each evaluation class by first listing any one component whose evaluation class is less than or equal to the evaluation class for the table. Then, add to the table all components that meet all of the following conditions:
 - a) They have an evaluation class less than or equal to the class of the table.
 - b) They receive data from the network at a level that is being sent by a component who is already in the table.
 - c) They send data to the network at a level that is equal to or less than any node already in the table.
4. After all the tables have been constructed then the Network Table Evaluation Class of each table is compared to the Maximum and Minimum for the Table with regard to the rules specified by the Environmental Guidelines.
5. If all Tables satisfy the assurance requirements for the Environmental Guidelines then the Network passes the assurance requirements. If any of the Tables provide a greater risk index than is permitted by the Environmental Guidelines then the Network provides a high level of risk, and should not be connected as currently designed.

Table C2. B2 TABLE 1

_____ — — — _____ —

(Note: Table not included)

Table C3. B2 TABLE 1, EXTENDED

(Note: Table not included)

Table C4:

Table C4(a). B2 TABLE 1

(Note: Table not included)

Table C4(b). B2 TABLE 2

(Note: Table not included)

C.4.1. Example B2 Table

As an example consider Table C2. This represents a B2 table under construction with a single entry. If in the network there existed another node which was evaluated at B2 and could receive and send at C-S, this node would be added to the table, producing Table C3. In contrast if there existed in the network a B2 node that could receive at S-TS but could only send at TS, this node would not be added to Table C3 but could be used to start a second B2 table. There would then be a set of two tables, represented in Table C4.

C.4.2. Sample Network and Tables

A sample network is illustrated in Figure C4. The tables that are produced for it are given in Tables C5(a) through C5(h).

Notice at the B2 level the network is represented by two tables, C5(b) and C5(c). This is due to the fact that one of the components (Component C) is a receive-only component. Such components will always end up in a table by themselves due to the fact that they never can affect the security of other nodes on the network. (The only security consideration to make with such components is whether they are receiving data at a level dominated by the approved maximum processing level for the component.)

Notice at the B1 level the components are each in a

table by themselves (Tables C5(d), C5(e), and C5(f)). This is due to the fact that although Component B may receive data from Component E, it never sends data to the network at a level lower than that sent by E (i.e., B can only send TS, never Confidential or Unclassified). Thus there is no "added" risk in having B receive data from E. If it were the case the B could send data at a level lower than (or equal to) E then they would be included in the same table since they present an added (or equal) risk.

C.5. Environmental Considerations

Because of the very nature of networks, it is necessary to say something about the assumptions made with respect to physical and logical protection of network links. While the requirements stated in Part II of this document apply to the single trusted system view, and not to the networks addressed by this Appendix, the issues are also important in the interconnected accredited AIS view. Therefore, this section presents a brief description of some of the important considerations. The interested reader is referred to Part II of this document for further detail.

It is not, repeat, NOT the intent of this Appendix to define the measures that are considered adequate under all circumstances. Rather, it is to identify the requirements, and where known, common methods of achieving the desired protection.

This Appendix establishes the requirement for integrity protection of control information in unclassified networks and the requirement to preserve the integrity of both control data and information in any other network.

C.5.1. Communications Integrity

The accreditor(s) will define transmission accuracy requirements for interconnecting two components. All elements involved in the interconnection of the components will demonstrate either by testing or by otherwise convincing argument that they meet the requirements. Since absolute transmission accuracy is not possible, a capability of testing for, detecting and reporting errors will be demonstrated. Acceptable methods of limiting errors include use of cryptographic checksums, protected wireways, and reliable delivery protocols used separately or in combination.

Hardware and/or software will be provided that can be used to periodically validate the correct operation of all elements involved in interconnecting two accredited components.

Trusted communications paths will be provided between network elements whenever secure element-to-element communications are required. (Initialization, cryptographic key management, change of subject or object security levels or access authorizations, etc.)

C.5.2. Denial of Service

The accreditor(s) will define denial of service conditions relative to the services being provided by the interconnected components. Hardware and or software will be provided to periodically assure the accessibility of the interconnected components.

C.5.3. Data Content Protection

Where classified information or sensitive but unclassified information is to be exchanged between logically connected components, it is the responsibility of each of the DAAs to assure that the content of their communication is

protected from unauthorized observation. Acceptable means of providing this protection include cryptography, and Protected Wireline Distribution Systems (PWDS).

Where the connection infrastructure is operated by a services organization for a number of different subscribers, suitable communications protection may be offered by the service organization. Regardless, it is the responsibility of the individual DAA to determine whether this protection is adequate for the information to be exchanged.

Figure C4. A Sample Network

(Note: Figure not included)

Component ID	Permitted Operations
A	Send and Receive data from C through TS
B	Send and Receive TS-only
C	Can Receive only audit records, all of which are treated as TS
D	Can Send and Receive C through TS
E	Can Send and Receive S-only
F	Send and Receive S and TS data

Table C5(a). NETWORK TABLE

NETWORK TABLE EVAL CLASS = A1
NETWORK TABLE MAXIMUM = TS
NETWORK TABLE MINIMUM = C
ENVIRONMENTAL GUIDELINES RULING = OK

(Note: Table not included)

(Note: since there are no B3 components, the B3 tables are identical to the B2 tables and are therefore not reproduced here.)

Table C5(b). B2 TABLE 1

TABLE EVAL CLASS = B2
TABLE MAXIMUM = TS
TABLE MINIMUM = S
ENV. GUIDELINES RULING= OK

Table C5(c). B2 TABLE 2

TABLE EVAL CLASS = B2
TABLE MAXIMUM = TS
TABLE MINIMUM = TS
ENV. GUIDELINES RULING= OK

(Note: Tables not included)

Table C5(d). B1 TABLE 1

TABLE EVAL CLASS = B1
TABLE MAXIMUM = S
TABLE MINIMUM = S
ENV. GUIDELINES RULING = OK

Table C5(e). B1 TABLE 2

TABLE EVAL CLASS = B1
TABLE MAXIMUM = TS
TABLE MINIMUM = TS
ENV. GUIDELINES RULING = OK

(Note: Table not included)

Acronyms

ACL	-	Access Control List
ADP	-	Automatic Data Processing
AIS	-	Automated Information System
ARPANET	-	Advanced Research Projects Agency Network
COMSEC	-	Communications Security
CPU	-	Central Processing Unit
CRC	-	Cyclic Redundancy Code or Cyclic Redundancy Check
DAA	-	Designated Approving Authority
DBMS	-	Data Base Management System
DAC	-	Discretionary Access Control
DDN	-	Defense Data Network
DoD	-	Department of Defense
DoDIIS	-	Department of Defense Intelligence Information System
DOS	-	Denial-of-service
DTLS	-	Descriptive Top-Level Specification
E3	-	End-to-end Encryption
FTLS	-	Formal Top-Level Specification
FTP	-	File Transfer Protocol
IP	-	Internet Protocol
I S/A AMPE Exchange	-	Inter-Service/Agency Automatic Message Processing
ISDN	-	Integrated Services Digital Network
ISO	-	International Standards Organization
KDC	-	Key Distribution Center
LAN	-	Local Area Network
LRC	-	Longitudinal Redundancy Check
MAC	-	Mandatory Access Control
MDC	-	Manipulation Detection Code
MSM	-	Message Stream Modification
MWT	-	Maximum Waiting Time
NSA	-	National Security Agency
NTCB	-	Network Trusted Computing Base
OSI	-	Open System Interconnection
PABX	-	Private Automated Branch Exchange
PDU	-	Protocol Data Unit (a.k.a. packet, datagram)
PKC	-	Public Key Cryptosystem
PWDS	-	Protected Wireline Distribution System
ROM	-	Read Only Memory
SACDIN	-	Strategic Air Command Digital Network
TAC	-	Terminal Access Controller
TCB	-	Trusted Computer Base
TCP	-	Transmission Control Protocol
TELNET	-	Network Virtual Terminal Protocol
TLS	-	Top Level Specification
TCSEC	-	Trusted Computer System Evaluation Criteria
TFE	-	Trusted Front-end Processor
TIU	-	Trusted Network Interface Unit

TNI - Trusted Network Interpretations
VMM - Virtual Machine Monitor

Glossary

- A -

-

Access - (1) A specific type of interaction between a subject and an object that results in the flow of information from one to the other. (2) The ability and the means necessary to approach, to store or retrieve data, to communicate with, or to make use of any resource of an ADP system.

Access control - (1) The limiting of rights or capabilities of a subject to communicate with other subjects, or to use functions or services in a computer system or network. (2) Restrictions controlling a subject's access to an object.

Access control list - (1) A list of subjects authorized for specific access to an object. (2) A list of entities, together with their access rights, which are authorized to have access to a resource.

Accountability - the quality or state which enables actions on an ADP system to be traced to individuals who may then be held responsible. These actions include violations and attempted violations of the security policy, as well as allowed actions.

Accreditation - the managerial authorization and approval, granted to an ADP system or network to process sensitive data in an operational environment, made on the basis of a certification by designated technical personnel of the extent to which design and implementation of the system meet pre-specified technical requirements, e.g., TCSEC, for achieving adequate data security. Management can accredit a system to operate at a higher/lower level than the risk level recommended (e.g., by the Requirements Guideline-) for the certification level of the system. If management accredits the system to operate at a higher level than is appropriate for the certification level, management is accepting the additional risk incurred.

Accreditation range - of a host with respect to a particular network, is a set of mandatory access control levels

- Security Requirements: Guidance for Applying the

for data storage, processing, and transmission. The accreditation range will generally reflect the sensitivity levels of data that the accreditation authority believes the host can reliably keep segregated with an acceptable level of risk in the context of the particular network for which the accreditation range is given. Thus, although a host system might be accredited to employ the mandatory access control levels CONFIDENTIAL, SECRET, and TOP SECRET in stand-alone operation, it might have an accreditation range consisting of the single value TOP SECRET for attachment to some network.

Audit trail - (1) A set of records that collectively provide documentary evidence of processing used to aid in tracing from original transactions forward to related records and reports, and/or backwards from records and reports to their component source transactions. (2) Information collected or used to facilitate a Security Audit.

Authentication - (1) To establish the validity of a claimed identity. (2) To provide protection against fraudulent transactions by establishing the validity of message, station, individual, or originator.

- B -

Bell-LaPadula Model - a formal state transition model of computer security policy that describes a set of access control rules. In this formal model, the entities in a computer system are divided into abstract sets of subjects and objects. The notion of a secure state is defined and it is proven that each state transition preserves security by moving from secure state to secure state; thus, inductively proving that the system is secure. A system state is defined to be "secure" if the only permitted access modes of subjects to objects are in accordance with a specific security policy. In order to determine whether or not a specific access mode is allowed, the clearance of a subject is compared to the classification of the object and a determination is made as to whether the subject is authorized for the specific access mode. The clearance/classifications scheme is expressed in terms of a lattice. See also: Lattice, Simple Security Property, *-Property. For further information see Bell, D. Elliott and LaPadula, Leonard J., Secure Computer Systems: Unified Exposition and MULTICS

- C -

Category - a grouping of objects to which a non-hierarchical restrictive label is applied (e.g., proprietary, compartmented information). Subjects must be privileged to access a category.

Certification - the technical evaluation of a system's security features, made as part of and in support of the approval/accreditation process, that establishes the extent to which a particular system's design and implementation meet a set of specified security requirements.

Closed user group - a closed user group permits users belonging to a group to communicate with each other, but precludes communications with other users who are not members of the group.

Communication channel - the physical media and devices which provide the means for transmitting information from one component of a network to (one or more) other components.

Communication link - the physical means of connecting one location to another for the purpose of transmitting and/or receiving data.

Compartment - a designation applied to a type of sensitive information, indicating the special handling procedures to be used for the information and the general class of people who may have access to the information. It can refer to the designation of information belonging to one or more categories.

Component - a device or set of devices, consisting of hardware, along with its firmware, and/or software that performs a specific function on a computer communications network. A component is a part of the larger system, and may itself consist of other components. Examples include modems, telecommunications controllers, message switches, technical control devices, host computers, gateways, communications subnets, etc.

Component Reference Monitor - an access control concept that refers to an abstract machine that mediates all access to objects within a component by subjects within the component.

Compromise - a violation of the security system such

that an unauthorized disclosure of sensitive information may have occurred.

Confidentiality - the property that information is not made available or disclosed to unauthorized individuals, entities, or processes.

Configuration control - management of changes made to a system's hardware, software, firmware, and documentation throughout the development and operational life of the system.

Connection - a liaison, in the sense of a network interrelationship, between two hosts for a period of time. The liaison is established (by an initiating host) for the purpose of information transfer (with the associated host); the period of time is the time required to carry out the intent of the liaison (e.g., transfer of a file, a chatter session, delivery of mail). In many cases, a connection (in the sense of this glossary) will coincide with a host-host connection (in a special technical sense) established via TCP or equivalent protocol. However a connection (liaison) can also exist when only a protocol such as IP is in use (IP has no concept of a connection that persists for a period of time). Hence, the notion of connection as used here is independent of the particular protocols in use during a liaison of two hosts.

Correctness - the extent to which a program satisfies its specifications.

Covert channel - a communications channel that allows a process to transfer information in a manner that violates the system's security policy. A covert channel typically communicates by exploiting a mechanism not intended to be used for communication. See Covert storage channel and Covert timing channel. Compare Overt channel.

Covert storage channel - a covert channel that involves the direct or indirect writing of a storage location by one process and the direct or indirect reading of the storage location by another process. Covert storage channels typically involve a finite resource (e.g., sectors on a disk) that is shared by two subjects at different security levels.

Covert timing channel - a covert channel in which one process signals information to another by modulating its own use of system resources (e.g., CPU time) in such a way that this manipulation affects the real response time observed by the second process.

Data confidentiality - the state that exists when data is held in confidence and is protected from unauthorized disclosure.

Data integrity - (1) The state that exists when computerized data is the same as that in the source documents and has not been exposed to accidental or malicious alteration or destruction. (2) The property that data has not been exposed to accidental or malicious alteration or destruction.

Dedicated Security Mode - the mode of operation in which the system is specifically and exclusively dedicated to and controlled for the processing of one particular type or classification of information, either for full-time operation or for a specific period of time. Compare Multilevel Security Mode, System High Security Mode.

Denial of service - the prevention of authorized access to system assets or services, or the delaying of time critical operations.

Descriptive top-level specification (DTLS) - a top-level specification that is written in a natural language (e.g., English), an informal program design notation, or a combination of the two.

Discretionary access control (DAC) - a means of restricting access to objects based on the identity of subjects and/or groups to which they belong. The controls are discretionary in the sense that: (a) A subject with a certain access permission is capable of passing that permission (perhaps indirectly) on to any other subject; (b) DAC is often employed to enforce need-to-know; (c) Access control may be changed by an authorized individual. Compare to Mandatory Access Control.

Domain - the set of objects that a subject has the ability to access.

Dominated by (the relation) - a security level A is dominated by security level B if the clearance/classification in A is less than or equal to the clearance/classification in B and the set of access approvals (e.g., compartment designators) in A is contained in (the set relation) the set of access approvals in B (i.e., each access approval appearing in A also appears in B). Depending upon the policy enforced (e.g., non-disclosure, integrity) the definition of "less than or equal to" and "contained in" may vary. For example, the level of an object of high integrity (i.e., an object which should be modifiable by very trustworthy individuals) may be defined to be "less than" the level of an object of low integrity (i.e., an object which is modifiable by everyone).

Dominates (the relation) - security level B dominates security level A if A is dominated by B.

- E -

Exploitable channel - any channel that is usable or detectable by subjects external to the Trusted Computing Base.

- F -

Flaw - an error of commission, omission, or oversight in a system that allows protection mechanisms to be bypassed.

Flaw hypothesis methodology - a system analysis and penetration technique where specifications and documentation for the system are analyzed and then flaws in the system are hypothesized. The list of hypothesized flaws is then prioritized on the basis of the estimated probability that a flaw actually exists and, assuming a flaw does exist, on the ease of exploiting it and on the extent of control or compromise it would provide. The prioritized list is used to direct the actual testing of the system.

Formal proof - a complete and convincing mathematical argument, presenting the full logical justification for each proof step, for the truth of a theorem or set of theorems. The formal verification process uses formal proofs to show the truth of certain properties of formal specification and for showing that computer programs satisfy their specifications. Automated tools may (but need not) be used to formulate and/or check the proof.

Formal security policy model - a mathematically precise statement of security policy. To be adequately precise, such a model must represent the initial state of a system, the way in which the system progresses from one state to another, and a definition of a "secure" state of the system. To be acceptable as a basis for a TCB, the model must be supported by a formal proof that if the initial state of the system satisfies the definition of a "secure" state and if all assumptions required by the model hold, then all future states of the system will be secure. Some formal modeling techniques include: state transition models, temporal logic models, denotational semantics models, algebraic specification models. See also: Bell-LaPadula Model, Security Policy Model.

Formal top-level specification (FTLS) - a Top-Level Specification that is written in a formal mathematical

language to allow theorems showing the correspondence of the system specification to its formal requirements to be hypothesized and formally proven.

Formal verification - the process of using formal proofs to demonstrate the consistency (design verification) between a formal specification of a system and a formal security policy model or (implementation verification) between the formal specification and its program implementation.

Functional testing - the portion of security testing in which the advertised features of a system are tested for correct operation.

- H -

Hierarchical decomposition - the ordered, structured reduction of a system or a component to primitives.

Host - any computer-based system connected to the network and containing the necessary protocol interpreter software to initiate network access and carry out information exchange across the communications network. This definition encompasses typical "mainframe" hosts, generic terminal support machines (e.g., ARPANET TAC, DoDIIS NTC), and workstations connected directly to the communications subnetwork and executing the intercomputer networking protocols. A terminal is not a host because it does not contain the protocol software needed to perform information exchange; a workstation (by definition) is a host because it does have such capability.

- I -

Integrity - See data integrity and integrity policy.

Integrity Policy - a security policy to prevent unauthorized users from modifying, viz., writing, sensitive information. See also Security Policy.

Internal subject - a subject which is not acting as direct surrogate for a user. A process which is not associated with any user but performs system-wide functions such as packet switching, line printer spooling, and so on. Also known as a daemon or a service machine.

- L -

Label - see Security Label and Sensitivity Label.

Lattice - a partially ordered set for which every pair of elements has a greatest lower bound and a least upper bound.

Least privilege - this principle requires that each subject in a system be granted the most restrictive set of privileges (or lowest clearance) needed for the performance of authorized tasks. The application of this principle limits the damage that can result from accident, error, or unauthorized use.

- M -

Mandatory access control - a means of restricting access to objects based on the sensitivity (as represented by a label) of the information contained in the objects and the formal authorization (i.e., clearance) of subjects to access information of such sensitivity.

Multilevel device - a device that is used in a manner that permits it to simultaneously process data of two or more security levels without risk of compromise. To accomplish this, sensitivity labels are normally stored on the same physical medium and in the same form (i.e., machine-readable or human-readable) as the data being processed.

Multilevel secure - a class of system containing information with different sensitivities that simultaneously permits access by users with different security clearances and needs-to-know, but prevents users from obtaining access to information for which they lack authorization.

Multilevel Security Mode - the mode of operation that allows two or more classification levels of information to be processed simultaneously within the same system when some users are not cleared for all levels of information present. Compare Dedicated Security Mode, System High Security Mode.

- N -

Network architecture - the set of layers and protocols (including formats and standards that different hardware/software must comply with to achieve stated objectives) which define a Network.

Network component - a network subsystem which is evaluatable for compliance with the trusted network interpretations, relative to that policy induced on the component by the overall network policy.

Network connection - A network connection is any logical or physical path from one host to another that makes

possible the transmission of information from one host to the other. An example is a TCP connection. But also, when a host transmits an IP datagram employing only the services of its "connectionless" Internet Protocol interpreter, there is considered to be a connection between the source and the destination hosts for this transaction.

Network Reference Monitor - an access control concept that refers to an abstract machine that mediates all access to objects within the network by subjects within the network.

Network security - the protection of networks and their services from unauthorized modification, destruction, or disclosure. Providing an assurance that the network performs its critical functions correctly and there are no harmful side-effects. Includes providing for information accuracy.

Network security architecture - a subset of network architecture specifically addressing security-relevant issues.

Network sponsor - the individual or organization that is responsible for stating the security policy enforced by the network, for designing the network security architecture to properly enforce that policy, and for ensuring that the network is implemented in such a way that the policy is enforced. For commercial, off-the-shelf systems, the network sponsor will normally be the vendor. For a fielded network system, the sponsor will normally be the project manager or system administrator.

Network System - a system which is implemented with a collection of interconnected network components. A network system is based on a coherent security architecture and design.

Network trusted computing base (NTCB) - the totality of protection mechanisms within a network system -- including hardware, firmware, and software -- the combination of which is responsible for enforcing a security policy. (See also Trusted Computing Base.)

NTCB Partition - the totality of mechanisms within a single network component for enforcing the network policy, as allocated to that component; the part of the NTCB within a single network component.

- 0 -

Object - a passive entity that contains or receives information. Access to an object potentially implies access to the information it contains. Examples of objects are:

records, blocks, pages, segments, files, directories, directory trees, and programs, as well as bits, bytes, words, fields, processors, video displays, keyboards, clocks, printers, etc. See also Passive.

Object reuse - the reassignment of a medium (e.g., page frame, disk sector, magnetic tape) that contained one or more objects to some subject. To be securely reassigned, such media must contain no residual data from the previously contained object(s).

OSI Architecture - the International Organization for Standardization (ISO) provides a framework for defining the communications process between systems. This framework includes a network architecture, consisting of seven layers. The architecture is referred to as the Open Systems Interconnection (OSI) model or Reference Model. Services and the protocols to implement them for the different layers of the model are defined by international standards. From a systems viewpoint, the bottom three layers support the components of the network necessary to transmit a message, the next three layers generally pertain to the characteristics of the communicating end systems, and the top layer supports the end users. The seven layers are: 1. Physical Layer: Includes the functions to activate, maintain, and deactivate the physical connection. It defines the functional and procedural characteristics of the interface to the physical circuit: the electrical and mechanical specifications are considered to be part of the medium itself. 2. Data Link Layer: Formats the messages. Covers synchronization and error control for the information transmitted over the physical link, regardless of the content. "Point-to point error checking" is one way to describe this layer. 3. Network Layer: Selects the appropriate facilities. Includes routing communications through network resources to the system where the communicating application is: segmentation and reassembly of data units (packets) ; and some error correction. 4. Transport Layer: Includes such functions as multiplexing several independent message streams over a single connection, and segmenting data into appropriately sized packets for processing by the Network Layer. Provides end-to-end control of data reliability. 5. Session Layer: Selects the type of service. Manages and synchronizes conversations between two application processes. Two main types of dialogue are provided: two-way simultaneous (full-duplex), or two-way alternating (half-duplex). Provides control functions similar to the control language in computer system. 6. Presentation Layer: Ensures that information is delivered in a form that the receiving system can understand and use. Communicating parties determine the format and language (syntax) of messages: translates if required, preserving the meaning (semantics). 7. Application Layer: Supports distributed applications by manipulating information. Provides resource management for file

transfer, virtual file and virtual terminal emulation, distributed processes and other applications.

Overt channel - an overt channel is a path within a network which is designed for the authorized transfer of data.

- P -

Passive - (1) A property of an object or network object that it lacks logical or computational capability and is unable to change the information it contains. (2) Those threats to the confidentiality of data which, if realized, would not result in any unauthorized change in the state of the intercommunicating systems (e.g., monitoring and/or recording of data).

Penetration - the successful violation of a protected system.

Penetration testing - the portion of security testing in which the penetrators attempt to circumvent the security features of a system. The penetrators may be assumed to use all system design and implementation documentation, which may include listings of system source code, manuals, and circuit diagrams. The penetrators work under no constraints other than those that would be applied to ordinary users or implementors of untrusted portions of the component.

Privacy - (1) the ability of an individual or organization to control the collection, storage, sharing, and dissemination of personal and organizational information. (2) The right to insist on adequate security of, and to define authorized users of, information or systems. Note: The concept of privacy cannot be very precise and its use should be avoided in specifications except as a means to require security, because privacy relates to "rights" that depend on legislation.

Process - a program in execution. It is completely characterized by a single current execution point (represented by the machine state) and address space.

Protection-critical portions of the TCB - those portions of the TCB whose normal function is to deal with the control of access between subjects and objects. See also Subject, Object, Trusted Computer Base.

Protection philosophy - an informal description of the overall design of a system that delineates each of the protection mechanisms employed. A combination (appropriate to the evaluation class) of formal and informal techniques is used to show that the mechanisms are adequate to enforce the

security policy.

- R -

Read - a fundamental operation that results only in the flow of information from an object to a subject.

Read access - permission to read information.

Reference monitor concept - an access control concept that refers to an abstract machine that mediates all accesses to objects by subjects. See also Security Kernel.

Reliability - the extent to which a system can be expected to perform its intended function with required precision.

Resource - anything used or consumed while performing a function. The categories of resources are: time, information, objects (information containers), or processors (the ability to use information). specific examples are: CPU time; terminal connect time; amount of directly-addressable memory; disk space; number of I/O requests per minute, etc.

- S -

Secrecy Policy - a security policy to prevent unauthorized users from reading sensitive information. See also Security Policy

Security architecture - the subset of computer architecture dealing with the security of the computer or network system. See computer architecture, network architecture.

Security-Compliant Channel - A channel is Security-Compliant if the enforcement of the network policy depends only upon characteristics of the channel either (1) included in the evaluation, or (2) assumed as a installation constraint and clearly documented in the Trusted Facility Manual.

Security Kernel - the hardware, firmware, and software elements of a Trusted Computing Base (or Network Trusted Computing Base partition) that implement the reference monitor concept. It must mediate all accesses, be protected from modification, and be verifiable as correct.

Security label - see Sensitivity label.

Security level - the combination of hierarchical clas-

sification and a set of non-hierarchical categories that represents the sensitivity of information.

Security policy - the set of laws, rules, and practices that regulate how an organization manages, protects, and distributes sensitive information.

Security policy model - an informal presentation of a formal security policy model.

Security testing - a process used to determine that the security features of a system are implemented as designed and that they are adequate for a proposed application environment. This process includes hands-on functional testing, penetration testing, and verification. See also: Functional Testing, Penetration Testing, Verification.

Sensitivity label - A piece of information that represents the security level of an object and that describes the sensitivity (e.g., classification) of the data in the object. Sensitivity labels are used by the NTCB as the basis for mandatory access control decisions.

Sensitivity level - See Security level.

Simple security property - a Bell-LaPadula security model rule allowing a subject read access to an object only if the security level of the subject dominates the security level of the object.

Single-level device - a device that is used to process data of a single security level at any one time. Since the device need not be trusted to separate data of different security levels, sensitivity labels do not have to be stored with the data being processed.

*-property (star property) - a Bell-LaPadula security model rule allowing a subject write access to an object only if the security level of the subject is dominated by the security level of the object. Also known as the Confinement Property.

Storage object - an object that supports both read and write accesses.

Subject - an active entity, generally in the form of a person, process, or device that causes information to flow among objects or changes the system state. Technically, a process/domain pair.

Subject security level - a subject's security level is equal to the security level of the objects to which it has both read and write access. A subject's security level must always be dominated by the clearance of the user the subject

is associated with.

System - an assembly of computer and/or communications hardware, software, and firmware configured for the purpose of classifying, sorting, calculating, computing, summarizing, transmitting and receiving, storing and retrieving data with the purpose of supporting users.

System High - the highest security level supported by a system at a particular time or in a particular environment.

System High Security Mode - the mode of operation in which system hardware and software is only trusted to provide discretionary protection between users. In this mode, the entire system, to include all components electrically and/or physically connected, must operate with security measures commensurate with the highest classification and sensitivity of the information being processed and/or stored. All system users in this environment must possess clearances and authorization for all information contained in the system. All system output must be clearly marked with the highest classification and all system caveats until the information has been reviewed manually by an authorized individual to ensure appropriate classifications and that caveats have been affixed. Compare Dedicated Security Mode, Multilevel Security Mode.

System Low - the lowest security level supported by a system at a particular time or in a particular environment.

System Security Officer (SSO) - the person responsible for the security of a system. The SSO is authorized to act in the "security administrator" role. Functions that the SSO is expected to perform include: auditing and changing security characteristics of a user.

- T -

Top-level specification (TLS) - a non-procedural description of system behavior at the most abstract level. Typically a functional specification that omits all implementation details.

Trap-door - a hidden software or hardware mechanism that permits system protection mechanisms to be circumvented. It is activated in some non-apparent manner (e.g., special "random" key sequence at a terminal).

Trojan horse - a computer program with an apparently or actually useful function that contains additional (hidden) functions that surreptitiously exploit the legitimate authorizations of the invoking process to the detriment of security. For example, making a "blind copy" of a sensitive

file for the creator of the Trojan Horse.

Trusted channel - a mechanism by which two NTCB partitions can communicate directly. This mechanism can be activated by either of the NTCB partitions, cannot be imitated by untrusted software, and maintains the integrity of information that is sent over it. A trusted channel may be needed for the correct operation of other security mechanisms.

Trusted computer system - a system that employs sufficient hardware and software integrity measures to allow its use for processing simultaneously a range of sensitive or classified information.

Trusted computing base (TCB) - the totality of protection mechanisms within a computer system -- including hardware, firmware, and software -- the combination of which is responsible for enforcing a security policy. It creates a basic protection environment and provides additional user services required for a trusted computer system. The ability of a trusted computing base to correctly enforce a security policy depends solely on the mechanisms within the TCB and on the correct input by system administrative personnel of parameters (e.g., a user's clearance) related to the security policy.

Trusted functionality - that which is determined to be correct with respect to some criteria, e.g. as established by a security policy. The functionality shall neither fall short of nor exceed the criteria.

Trusted path - a mechanism by which a person at a terminal can communicate directly with the Trusted Computing Base. This mechanism can only be activated by the person or the Trusted Computing Base and cannot be imitated by untrusted software.

Trusted software - the software portion of a Trusted Computing Base.

Trusted subject - a subject that is part of the TCB. It has the ability to violate the security policy, but is trusted not to actually do so. For example in the Bell-LaPadulla model a trusted subject is not constrained by the *-property and thus has the ability to write sensitive information into an object whose level is not dominated by the (maximum) level of the subject, but it is trusted to only write information into objects with a label appropriate for the actual level of the information.

- U -

User - any person who interacts directly with a network system. This includes both those persons who are authorized to interact with the system and those people who interact without authorization (e.g., active or passive wiretappers). Note that "users" does not include "operators," "system programmers," "technical control officers," "system security officers," and other system support personnel. They are distinct from users and are subject to the Trusted Facility Manual and the System Architecture requirements. Such individuals may change the system parameters of the network system, for example by defining membership of a group. These individuals may also have the separate role of users.

- V -

Verification - the process of comparing two levels of system specification for proper correspondence (e.g., security policy model with top-level specification (TLS), TLS with source code, or source code with object code). This process may or may not be automated.

Virus - malicious software, a form of Trojan horse, which reproduces itself in other executable code.

- W -

Write - a fundamental operation that results only in the flow of information from a subject to an object.

Write access - permission to write an object.

References

Abrams, M. D. and H. J. Podell, Tutorial: Computer and Network Security, IEEE Computer Society Press, 1987.

Addendum to the Transport Layer Protocol Definition for Providing Connection Oriented End-to-End Cryptographic Data Protection Using A 64-bit Block Cipher, ISO TC 97 / SC 20 / WG 3, N 37, January 10, 1986.

Biba K. J., Integrity Considerations for Secure Computer Systems, MTR-3153, The MITRE Corporation, June 1975; ESD-TR-76-372, April 1977.

Bell, D. Elliot and LaPadula, Leonard J., Secure Computer Systems: Unified Exposition and Multics Interpretation, MTR 2997, The MITRE Corporation, April 1974. (AD/A 020 445)

Denning, D .E., Lunt, T. F., Neumann, P. G., Schell, R. R., Heckman, M. and Shockley, W., Secure Distributed Data Views, Security Policy and Interpretation for a Class A1 Multilevel Secure Relational Database System, SRI International, November 1986.

Girling, C. G., "Covert Channels in LAN's," IEEE Transactions on Software Engineering, Vol. SE-13, No. 2, February 1987.

Grohn, M. J., A Model of a Protected Data Management System, ESD-TR-76-289, I. P. Sharp Assoc. Ltd., June, 1976.

``Integrity and Inference Group Report,`` Proceedings of the National Computer Security Center Invitational Workshop on Database Security, Baltimore, MD, 17-20 June 1986.

ISO 7498/Part 2 - Security Architecture, ISO / TC 97 / SC 21 / N1528 / WG 1 Ad hoc group on Security, Project 97.21.18, September 1986.

Jueneman, R. R, "Electronic Document Authentication," IEEE Network Magazine, April 1987, pp 17-23.

Lipner, Steven B., ``Non-Discretionary Controls for Commercial Applications``, IEEE Proceedings of the 1982 Symposium on Security and Privacy, April 26-28, 1982, Oakland, CA.

National Computer Security Center, Department of Defense Password Management Guideline, CSC-STD-002-85, 12 April 1985.

National Computer Security Center, Department of Defense Trusted Computer Security Evaluation Criteria, DOD 5200.28-STD, December 1985.

National Computer Security Center, Security Requirements: Guidance for Applying the Department of Defense Trusted Computer System Evaluation Criteria in Specific Environments, CSC-STD-003-85, 25 June 1985.

Padlipsky, M. A., Snow, D. P., and Karger, P. A., Limitations of End-to-End Encryption in Secure Computer Networks, The MITRE Corporation, MTR-3592, Vol. I, May 1978 (ESD TR 78-158, DTIC AD A059221).

The Directory - Authentication Framework (Melbourne, April 1986), ISO/CCITT Directory Convergence Document #3.

Voydock, Victor L. and Stephen T. Kent, "Security Mechanisms in High-Level Network Protocols," Computing Surveys, Vol. 15, No. 2, June 1983, pp 135-171.

ISO developmental documents are of limited lifetime and availability.